



YOUR PC... INSIDE-OUT

TROUBLESHOOTING & MAINTAINING YOUR IBM PC & COMPATIBLES

ALSO FEATURING THE PS/2

HANDBOOK

 **DATA-TECH
INSTITUTE**

YOUR PC...INSIDE - OUT

**TROUBLESHOOTING AND MAINTAINING YOUR IBM PC
AND COMPATIBLES**

ALSO FEATURING THE PS/2

HANDBOOK

PROGRAM DEVELOPER AND WRITER

Mark Minasi
Moulton, Minasi and Company
7146 Rivers Edge Road
Columbia, MD 21044

PRODUCED AND PUBLISHED BY

Data-Tech Institute
Lakeview Plaza
57 Lakeview Avenue
Clifton, NJ 07015

Your PC...Inside-Out Videotape Package

Copyright (c) 1987 Data-Tech Institute and Mark Minasi
All Rights Reserved

Do Not Copy Without Written Authorization

Exclusive Marketing and Distribution: Data-Tech Institute
Lakeview Plaza
P.O. Box 2429
Clifton, NJ 07015

Printed in U.S.A.

IMPORTANT NOTE

A few words of disclaimer:

1. This text should be used in conjunction with the videotape.
2. Many products are mentioned in this text. We are not endorsing these products. Where it is noted, they have been of value. However, manufacturer quality can vary, and goods can be redesigned.
3. Please be sure not to try anything on the PC before you know what you are doing.
4. All programs on the Program Disk are public domain or shareware: we are NOT selling these programs, nor do we offer any license or guarantees about their performance. Please use them with caution.
5. Read all manuals which accompany your computer before beginning any maintenance or troubleshooting.
6. While every precaution has been taken in the preparation of this book, the publisher and writer assume no responsibility for errors or omissions. Neither is any liability assumed for damages resulting from the use of the information contained herein.

WARNING: Check your manufacturer's warranty. Please be aware that opening or otherwise modifying your PC (including the IBM PC) may void any manufacturer's warranty on the product.

YOUR PC...INSIDE - OUT

VIDEOTAPE

TABLE OF CONTENTS

The enclosed two (2) hour Videotape contains a tremendous amount of information and is coded for Level 1 Interactive Access enabling you to make full use of the pause, fast scan and rewind features on your VCR. This provides the opportunity to learn at your own pace.
Refer to the color codes in this section.

	COLOR
THE PC FAMILY	Blue
THE TOOLS	Pink
INTRODUCTION OF PC COMPONENTS	Red
DISASSEMBLING & ASSEMBLING THE PC	Grey
1. Display Boards	
2. The Hard Disk Controller	
3. The Floppy Drives	
4. The Hard Disk Drive	
5. The Power Supply	
6. The System Board	
DISASSEMBLING THE PS/2 MODEL 50	Lt. Blue
TROUBLESHOOTING APPROACHES	Green
1. Removing & Installing Semiconductor Chips	
2. Identifying Chips	
3. Chip Handling	
4. Chip Tips	
5. Configuration & Installation	
6. Configuration With PS/2s	

TROUBLESHOOTING AND MAINTAINING YOUR IBM PC AND COMPATIBLES

ALSO FEATURING THE PS/2

HANDBOOK

TABLE OF CONTENTS

1.	INTRODUCTION	1
1.1.	THE PURPOSE OF THIS PROGRAM	1
1.2.	STRUCTURE OF THIS PROGRAM.....	1
2.	INSIDE THE PC: PIECES OF THE PICTURE	2
2.1.	MODULARITY	2
2.2.	THE SYSTEM BOARD / MOTHERBOARD	2
2.3.	CENTRAL PROCESSING UNIT (CPU).....	2
2.4.	PLANAR MEMORY	5
2.5.	READ ONLY MEMORY (ROM)	5
2.6.	EMS, LIM, EEMS, PAGED, EXPANDED MEMORY	6
2.7.	THE PC BUS, THE MCA BUS, AND EXPANSION SLOTS	6
2.8.	SYSTEM CLOCK	7
2.9.	NUMERIC COPROCESSOR.....	7
2.10.	POWER SUPPLY	7
2.11.	KEYBOARD	8
2.12.	DISPLAY AND DISPLAY ADAPTERS	8
2.13.	FLOPPY AND HARD DISK CONTROLLER AND DISK DRIVES.....	8
2.14.	PRINTERS AND PLOTTERS.....	9
2.15.	SYSTEM CLOCK/CALENDAR.....	9
2.16.	MODEMS AND COMMUNICATION PORTS.....	9

3.	AVOIDING SERVICE: PREVENTATIVE MAINTENANCE.....	10
3.1.	INTRODUCTION	10
3.2.	STEP ONE TO MAINTENANCE: HOW DO I GET INSIDE THE PC?... 10	
3.2.1.	PC Disassembly Steps.....	10
3.3.	AVOIDING DUST	10
3.4.	MAGNETISM.....	11
3.5.	ELECTROMAGNETIC NOISE	12
3.5.1.	Electromagnetic Interference.....	12
3.5.2.	Power-Up Power Surges	12
3.5.3.	Static.....	13
3.5.4.	Crosstalk.....	14
3.5.5.	Radio Frequency Interference	15
3.6.	AVOIDING WATER AND LIQUIDS	16
3.6.1.	Corrosion	17
3.7.	HEAT.....	17
3.7.1.	Duty Cycles.....	18
3.7.2.	Thermal Shock	18
3.7.3.	Other Heat Effects.....	18
4.	TROUBLESHOOTING PROBLEMS:	
	WHAT TO DO WHEN SOMETHING GOES WRONG.....	20
4.1.	GENERAL TROUBLESHOOTING RULES	20
4.2.	TROUBLESHOOTING STEPS	20
4.3.	CHECK FOR OPERATOR ERROR	21
4.4.	IS EVERYTHING PLUGGED IN?.....	21
4.5.	CHECK THE SOFTWARE	22
4.5.1.	Resource Conflicts Among Software Packages	22
4.5.2.	Poorly Terminated Software.....	22
4.5.3.	Hardware-Related Software Faults	23
4.5.4.	Faulty Software	23
4.6.	WHAT AM I DOING DIFFERENTLY?.....	23
4.7.	CHECK EXTERNAL SIGNS.....	24
4.8.	RUN DIAGNOSTIC PROGRAMS	24
4.8.1.	Running the IBM Diagnostic Programs	25
4.8.2.	The Power On Self Test (POST).....	25
4.8.3.	IBM Diagnostic Error Codes	26
5.	CIRCUIT BOARDS AND CHIPS.....	30
5.1.	BURN-IN.....	30
5.2.	BOARD IDENTIFICATION: BOARDS FOUND IN THE PC.....	30
5.2.1.	Maintenance Considerations for the PS/2 Family	31
5.3.	INSTALLING NEW CIRCUIT BOARDS.....	31
5.3.1.	Circuit Board Removal and Insertion.....	31
5.3.2.	Installing and Configuring Boards: DMA, IRQ, and I/O Ports.....	32
5.3.3.	I/O Ports.....	33
5.3.4.	DMA Channels and IRQ Lines	33
5.3.5.	ROM Addresses.....	35
5.3.6.	A Hint For Connecting Cables and Edge Connectors	36
5.3.7.	Power Requirements.....	36
5.3.8.	A Configuration Example	36
5.3.9.	PS/2 Configuration Notes.....	37
5.4.	PROBLEMS WITH EXISTING BOARDS:	
	COMMON FAILURE POINTS	37

5.5.	FINDING AND REPLACING BAD BOARDS	38
5.5.1.	Identifying The Problem Board: Symptoms and Possibilities	38
5.6.	FINDING AND REPLACING BAD CHIPS	40
5.6.1.	Identifying Chips.....	41
5.6.2.	Limitations: How Far Should I Go In Isolating Bad Chips?.....	42
5.6.3.	Locating Chips With Heat Problems: Intermittent Failures.....	42
5.6.4.	Soldering and Desoldering.....	43
5.6.5.	Chip Sockets and Chip Insertion/Removal.....	44
5.7.	CPU CHIPS.....	44
5.7.1.	Common PC CPUs.....	45
5.7.2.	The 8088.....	45
5.7.3.	The 80286.....	45
5.7.4.	CPU Speeds.....	45
5.7.5.	Word Size.....	46
5.7.6.	Data Path.....	46
5.7.7.	Memory Size.....	46
5.8.	TROUBLESHOOTING THE SYSTEM BOARD	46
5.8.1.	Configuration With SETUP and DIP Switches	46
5.8.2.	PC System Board DIP Switch Settings	47
5.8.3.	Using the AT SETUP Program	47
5.8.4.	Test Points for The System Board	50
5.9.	SEMICONDUCTOR MEMORY.....	52
5.9.1.	Reading Memory Error Messages.....	53
5.9.2.	Reading Parity Check Messages	53
5.9.3.	Other Causes of Memory Errors	54
5.9.4.	Handling Parity Check A Little More Gently: NMI.COM.....	58
5.9.5.	Tips on Installing Memory Chips	58
6.	POWER SUPPLIES AND POWER PROTECTION.....	59
6.1.	COMPONENTS OF THE POWER SUPPLY	59
6.2.	MAINTENANCE.....	59
6.3.	UPGRADING THE POWER SUPPLY	59
6.4.	TROUBLESHOOTING THE POWER SUPPLY	60
6.4.1.	The Power Supply Troubleshooting Trail.....	61
6.4.2.	Replacing a Power Supply.....	61
6.5.	PROTECTING THE POWER SUPPLY	61
6.5.1.	Surge Suppressors and Spike Isolators	61
6.5.2.	Backup Power Supplies.....	63
7.	HARD DISK DRIVE MAINTENANCE AND FAILURE RECOVERY	64
7.1.	DISK TERMINOLOGY	64
7.1.1.	Cylinders, Heads, Platters, Tracks and Sectors.....	64
7.1.2.	Disk Performance I: Seek Times, and Latency Periods.....	69
7.1.3.	Disk Performance II: Data Transfer Rates and Interleave Factors	69
7.1.4.	Testing Disk Performance: CORETEST.....	74
7.1.5.	How Much Faster Is A Hard Disk Than a Floppy Disk?.....	75
7.2.	COMPONENTS: THE CONTROLLER AND THE DRIVE	75
7.3.	THE CONTROLLER.....	75
7.3.1.	The XT Controller vs the AT Controller	75
7.3.2.	MFM and RLL Controllers	76
7.4.	INTERFACES	77
7.4.1.	Informing the Controller of Drive Type	77
7.5.	THE DRIVE: THE "WINCHESTER" SEALED DRIVE	77
7.5.1.	Hard Cards.....	78

7.5.2.	You Can't Fix It, So You Must Protect It.....	78
7.5.3.	Band Steppers and Voice Coils.....	78
7.5.4.	Squeaky Drives.....	79
7.6.	THE BEST INSURANCE: GOOD BACKUPS	79
7.6.1.	Software Based Backup.....	80
7.6.2.	Backup Hardware	81
7.7.	DRIVE PREPARATION.....	81
7.7.1.	Physical Formatting.....	81
7.7.2.	Partitioning.....	82
7.7.3.	DOS Formatting.....	83
7.8.	FAILURE RECOVERY.....	83
7.8.1.	Damaged Boot Record.....	83
7.8.2.	Recovering From a Damaged FAT.....	85
7.8.3.	Recovery Lost Data Due to Disk Media Failure	86
7.8.4.	Recovering Accidentally Erased Files	87
7.8.5.	Recovering Accidentally Formatted Hard Disks.....	88
7.9.	MAINTENANCE: DISK OPTIMIZATION	88
7.9.1.	Fragmentation of Files.....	88
7.9.2.	Ever-Growing Subdirectories	89
7.9.3.	Disk Caching Programs	89

8.	FLOPPY DISK DRIVES.....	91
8.1.	FLOPPIES AND THE PC.....	91
8.2.	PIECES OF THE PICTURE.....	91
8.2.1.	The Floppy Diskette	91
8.2.2.	The Floppy Drive	92
8.2.3.	The Disk Controller Board.....	93
8.2.4.	The Cable	93
8.3.	MAINTENANCE.....	94
8.3.1.	How Often Should You Clean The Heads?.....	94
8.3.2.	Environmental Factors to Protect Disks and Diskettes.....	94
8.3.3.	A Disk Drive Tester.....	95
8.4.	PROBLEMS AND DIAGNOSIS	95
8.4.1.	Software Possibilities.....	95
8.4.2.	Hardware Possibilities.....	95
8.4.3.	Drive Head is Dirty	95
8.4.4.	Diskette is Bad.....	96
8.4.5.	Cable Is Loose.....	96
8.4.6.	Cable Is Bad	96
8.4.7.	Controller Is Bad.....	96
8.4.8.	Power Supply To Floppy Malfunctioning	96
8.4.9.	Drive is Malfunctioning	96
8.5.	COMMON REPAIR POINTS.....	97
8.5.1.	Is It Worth Repairing Floppy Drives?	97
8.5.2.	Removing Floppy Drives	98
8.5.3.	Installing a New Floppy.....	98
8.5.4.	More on Drive Select	99
8.5.5.	Replacing The Drive Door (Full Height)	99
8.5.6.	The Floppy Drive Electronics.....	99
8.5.7.	Saving A Dirty Floppy	100
8.6.	ADJUSTING FLOPPIES.....	100
8.6.1.	Disk Rotation Speed	100
8.6.2.	Testing And Adjusting Disk Speed	101
8.6.3.	Head Alignment	102

8.6.4.	Aligning the Heads With a Test Program	102
9.	PRINTERS AND PRINTER INTERFACES.....	104
9.1.	COMPONENTS	104
9.2.	MAINTENANCE.....	104
9.2.1.	Daisy Wheel Printers	104
9.2.2.	Dot Matrix Printers.....	105
9.2.3.	Laser Printers.....	105
9.3.	COMMON PROBLEMS AND TROUBLESHOOTING	106
9.3.1.	Isolate The Problem.....	106
9.3.2.	Cable Lengths	106
9.3.3.	Setup DIP Switch Problems	107
9.3.4.	Port Problems	107
9.3.5.	Software Problems.....	108
9.3.6.	The Mysterious Timeout	108
9.3.7.	The Weather	109
10.	MODEMS AND SERIAL INTERFACES	110
10.1.	INTRODUCTION	110
10.1.1.	Asynchronous Communication.....	110
10.1.2.	Synchronous Communication	110
10.1.3.	Local Area Networks (LANs)	111
10.2.	COMPONENTS	111
10.2.1.	The Asynchronous Port	111
10.2.2.	The Cable	111
10.2.3.	The Modem.....	112
10.2.4.	The Communications Software.....	112
10.3.	MAINTENANCE.....	112
10.4.	TROUBLESHOOTING.....	113
10.4.1.	The Software First	113
10.4.2.	The Port	114
10.4.3.	Testing The Port.....	115
10.4.4.	The Cable	115
10.4.5.	The Modem.....	116
10.4.6.	The Phone Line.....	116
10.4.7.	The Other Side.....	116
10.4.8.	Going Further	116
10.5.	COMMON CABLES	117
10.5.1.	25 Pin PC to 25 Pin Modem.....	118
10.5.2.	9 Pin AT to 25 Pin Modem	119
10.5.3.	PC to PC Null Modem	120
10.5.4.	PC to AT Null Modem	121
10.5.5.	AT to AT Null Modem	122
11.	KEYBOARDS	123
11.1.	DESIGN AND COMPONENTS.....	123
11.1.1.	The System Board Keyboard Interface.....	123
11.1.2.	The Keyboard Connector	124
11.2.	MAINTENANCE.....	124
11.3.	TROUBLESHOOTING.....	124
11.3.1.	Is It Plugged In?.....	124
11.3.2.	Is It One Key, Or All Keys?	124
11.3.3.	Checking the Cable Continuity	124
11.3.4.	Completely Disassembling the Keyboard	125

11.4.	REPLACEMENT KEYBOARDS.....	125
12.	DISPLAYS AND DISPLAY ADAPTERS.....	126
12.1.	COMPONENT: DISPLAYS.....	126
12.2.	COMPONENT: DISPLAY ADAPTER BOARDS	127
12.3.	UNDERSTANDING THE TERMINOLOGY	127
12.3.1.	Interlacing.....	128
12.4.	MAINTENANCE.....	128
12.5.	TROUBLESHOOTING.....	129
13.	THE MOST COMMON PROBLEM: SOFTWARE.....	130
13.1.	TRACKING DOWN THE PROBLEM.....	130
13.1.1.	If The Software Has Worked Before	130
13.1.2.	If You Are Installing Software Now.....	131
13.2.	REPORTING THE PROBLEM.....	131
14.	ADVANCED PROJECTS: SAFELY IMPROVING YOUR PC OR AT.....	133
14.1.	SPEAKER VOLUME CONTROL.....	133
14.2.	AN INTERFACE FOR AN EXTERNAL SPEAKER	133
14.3.	INSTALLING A HARDWARE RESET SWITCH	133
15.	APPENDICES.....	134
15.1.	THE BASIC SURVIVAL KIT.....	134
15.2.	THE PROGRAM DISK (PUBLIC DOMAIN UTILITIES).....	134

1. INTRODUCTION

1.1. THE PURPOSE OF THIS PROGRAM

You go to turn on your PC. It emits a long and a short beep, then refuses to do any more. You turn it on and off again, and it does the same thing. What now? Call the serviceman? The local PC expert? Open the yellow pages?

Corporations and private individuals spend millions of dollars each year for service contracts and repair bills for their PCs. Maintenance can be performed easily by anyone for a savings of *time* and *money*.

Some of these services are quite expensive — for instance, one maintenance company charges \$130/year for a service contract on a disk drive whose purchase price is around \$90!

You can save time, also. If you call the serviceman, how long will it take to get fixed? (Several weeks is not uncommon.) If you know how to do it yourself, you can fix many problems in under an hour.

This program won't teach you to fix all problems. Not all problems can be fixed — for instance, leaving your hard disk out in the rain will probably render your erstwhile data storage medium usable for little more than a paperweight. Nevertheless, even if you've never opened up a PC or installed an expansion board, this program will help you.

Many problems arise when installing new equipment, and installation is a related topic, anyway, so a lot of this material looks at installation.

In the process of working with IBM PCs since 1981 and microcomputers since 1976, I have picked up a number of very useful utilities which assist in the diagnosis of PC problems. These programs and their uses will be discussed in this text. Public domain programs are included on the program disk.

By the way, a note on terminology: when I say "PC" in this text, I include the XT and the AT, unless otherwise specified. The XT is virtually identical to the PC, except it has (1) a larger power supply, (2) three more expansion slots, and (3) no cassette port (how sad). The AT, on the other hand, is a radically different machine under the hood.

My goal in developing this program is to include material of use to "techies" as well as to those who have never even opened up a PC. I'm not going to try to make an electrical engineer out of you: I'm not one, myself. All that it takes to do most PC maintenance is a screwdriver and some patience.

1.2. STRUCTURE OF THIS PROGRAM

In the material that follows, I focus first on general preventative maintenance, then on how to figure out what's wrong when problems arise. Finally, separate chapters examine components, maintenance, common troubleshooting techniques and adjustments.

2. INSIDE THE PC: PIECES OF THE PICTURE

Never been inside your PC? Step inside...

2.1. MODULARITY

The PC is a modular device. This modularity makes problem determination and repair much more tractable than, say, repairing your TV. (Also, your PC lacks the large *capacitors* that make the TV dangerous to fool around with even when unplugged.) The vast majority of PC problems can be solved by simply isolating the faulty module and replacing it. Therefore, first you must know what modules make up a PC.

A PC is composed of just a few components:

System Board or Motherboard, containing:

CPU (Central Processing Unit)

Bus

Expansion Slots

Planar Memory

System Clock

Numeric Coprocessor

Keyboard Adapter (Interface)

Power Supply

Keyboard

Display and Display Adapter

Floppy Disk Controller and Floppy Disk Drives

Hard Disk Controller and Hard Disk Drive(s)

Multifunction Board, containing:

Printer and Printer Port

System Clock/Calendar

Modem and Serial Port

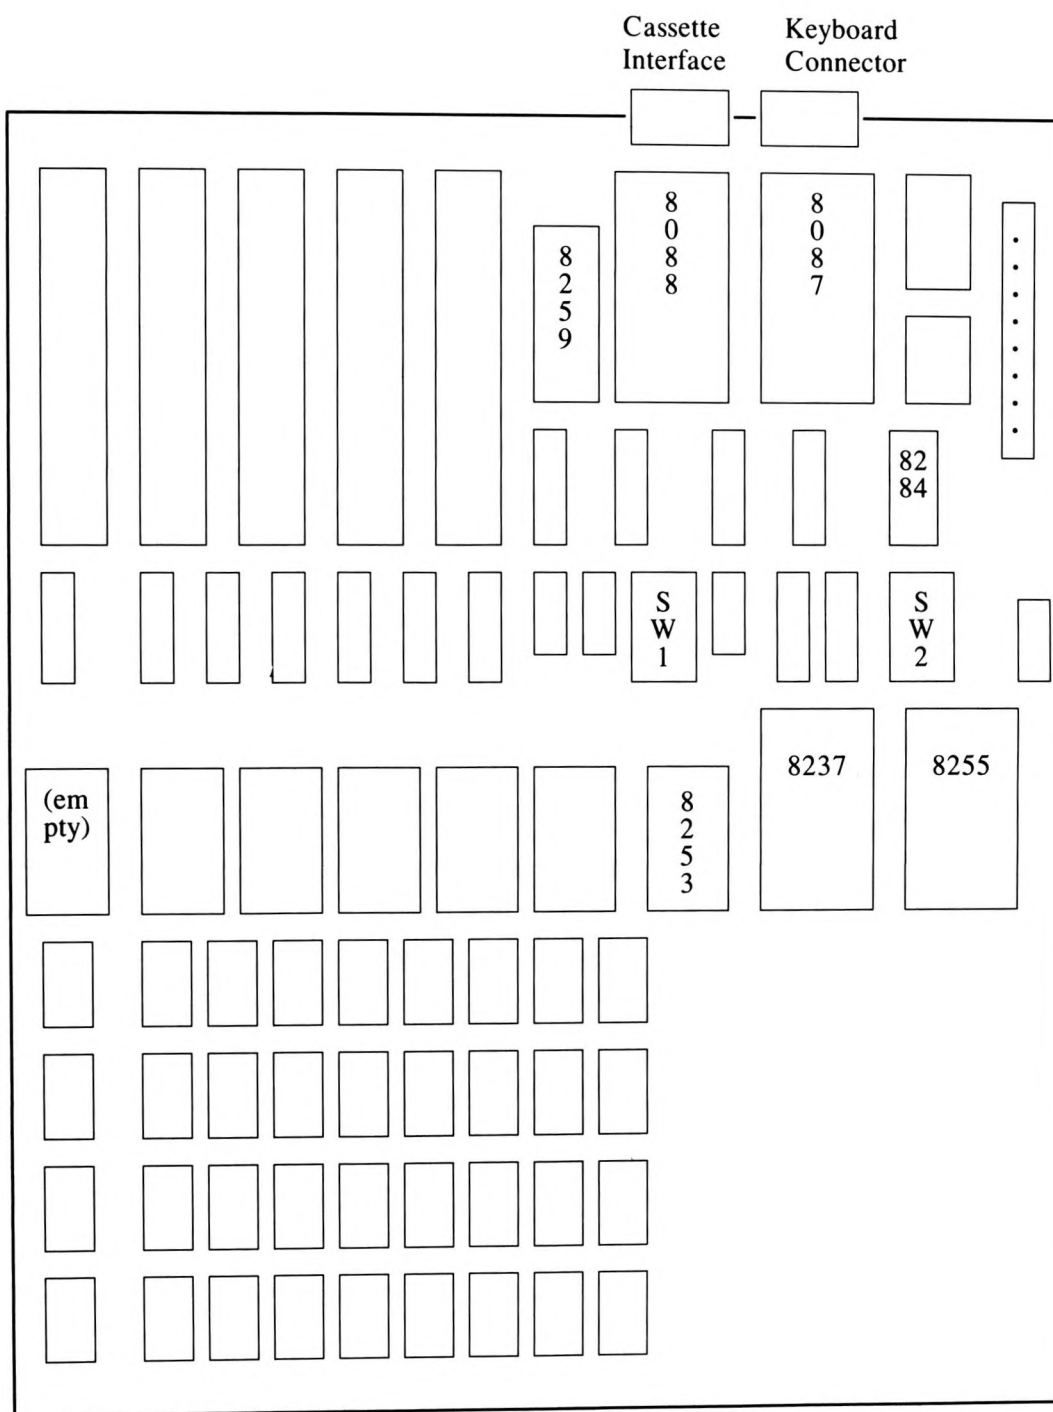
2.2. THE SYSTEM BOARD / MOTHERBOARD

Since their creation in 1974, microcomputers have usually included most of their essential electronics on a single printed circuit board, called the *motherboard*. The IBM PC and compatibles are no exception (with the sole exception of the Kaypro PC), but IBM chose to use a different term. Rather than calling this main circuit board a motherboard, Big Blue called it the *system board*. I will use both terms interchangeably in this text.

The newer PS/2 machines have yet another name for the motherboard: the *planar* board.

2.3. CENTRAL PROCESSING UNIT (CPU)

The "heart" of the PC is a microprocessor chip created by Intel Corporation or some other company licensed by Intel. Intel has designed and created many microprocessors over the years, but the ones that interest us are the 8088 and the 80286.



THE IBM PC SYSTEM BOARD

The PC Family Members

I. The PC/XT Family

Characteristics: 8 bit bus, 8088/8086 processors

Example: PC (5 slots)
 XT (8 slots)
 Convertible
 PS/2 Model 30
 Clones

II. The PC/AT Family

Characteristics: 16 bit bus (8 bit PC bus extended), 80286 processor

Examples: AT
 Deskpro 286, Compaq III
 XT 286
 Zenith 248
 Sperry IT

III. PS/2 Model 50/60

Characteristics: 16 bit Micro Channel Bus, 80286

Examples: Model 50, Model 60

Hardware incompatible with AT

Software mainly compatible with AT

IV. PS/2 Model 80

Characteristics: 32 bit Micro Channel Bus, 80386

Example: Model 80

V. Compaq 386

Characteristics: AT bus, 1 Compaq 32 bit slot for memory. Like an AT, but with an 80386 processor. Faster than an AT, but hardware/software compatible.

Examples:

 Compaq Deskpro 386
 CCI 386
 Trillian Power Systems PS/386
 ALR Access 386

CPU performance determines, in part, computer performance. There are a few CPU characteristics that will be examined at various points in this text. They are described in the following table:

Characteristic	Description	Unit	Typical Range
CPU Speed	How many operations can be done per second?	Mhz	4.77 Mhz - 20 Mhz
Word Size	What is the largest number that can be operated on in one operation?	bits	8 bits - 32 bits
Data Path	What is the largest number that can be transported into the chip in one operation?	bits	8 bits - 32 bits
Max Memory	How much memory can the chip use?	MB	1 MB - 2,000 MB

Table: CPU Characteristics

2.4. PLANAR MEMORY

The PC, like all computers, must have *main memory*. Main memory is high speed memory that the CPU can read from or write to. "High speed" here means less than a microsecond. The other name for such memory is RAM, Random Access Memory.

The PC can use 640K of main memory. The original PC's had an old system board, called the "PC-1" system board. The PC-1 motherboard could only accommodate 4 banks of 16K memory chips, so only 64K memory could be put on the motherboard. In 1983, IBM changed the PC motherboard so that (among other things) 4 rows of 64K memory could be put on it, for a total of 256K. The third recent motherboard was introduced in 1986, allowing two banks of 256K chips (512K total) and two banks of 64K chips (128K total) for a grand motherboard total of 640K. Virtually all clones available since 1985 have allowed 640K on the motherboard. Of course, the most recent motherboard is on the new (as of April 1987) IBM Model 30, a faster new XT-type machine. The old crop of PC and AT designs will be around for quite a while, so staying with them won't be much of a handicap.

For those millions of PCs without the room on the motherboard for 640K, supplemental memory had to come from expansion boards like Quadram's Quadboard or AST's SixPakPlus.

The term for main memory on the PC system board is *planar memory*.

2.5. READ ONLY MEMORY (ROM)

Another kind of memory exists, memory which cannot be altered. It can only be read, so it is called Read Only Memory. ROMs are found on expansion boards like EMS, LAN, or EGA cards. It is also found on the system board. The ROM on the system board is an important one called the BIOS ROM. The BIOS determines in large measure how compatible your PC compatible is.

ROM contains software, and software changes from time to time. Occasionally, a problem can be fixed by "upgrading the ROM" — getting the latest version of the ROM-based software from the manufacturer. Probably the worst offender here are 3278 terminal emulation boards like the IRMA board from Digital Communications Associates. (The new IRMA 2 does not use ROM, but rather loads its software from disk.)

For this reason, you've got to know exactly what ROMs are in the computers that you are responsible for. In a maintenance notebook, keep track of the serial numbers or dates on the labels pasted on the backs of the ROMs in your PCs. Whenever you install a board, note any ROM identifying marks. It will save you having to pop the top to find out when you call for service.

2.6. EMS, LIM, EEMS, PAGED, EXPANDED MEMORY

This does not appear on the system board, but as long as we're talking about memory this is a good time to discuss so-called "expanded" memory.

As everyone knows, 640K of memory — the most DOS will allow — is insufficient for most needs. Spreadsheets are a very popular application for the PC, but they have the flaw that the entirety of their data must be in main memory. Thus, a data file of a megabyte or two is impossible for a spreadsheet under DOS.

Lotus, Intel, and Microsoft (LIM) got together and developed a standard for a product that bypasses the DOS limitations through *memory paging*. Up to 8 megabytes of paged memory can be installed in a PC. This is not a panacea: only software written specially for the paged (called "expanded") memory can use this memory. There are only a few applications that do this: Lotus is the best known. It is sometimes known as the Expanded Memory Standard (EMS).

Repeating: *expanded* memory is Lotus/Intel/Microsoft paged memory. PCs, XTs, and ATs can use it. Another kind of memory exists for ATs. ATs can have up to 16 MB of main memory. (PCs and XTs cannot accommodate more than 640K.) This extra memory is called *extended* memory. Under DOS, this is useless memory: if you put it in, it won't do anything for you. Under OS/2, it will be usable.

2.7. THE PC BUS, THE MCA BUS, AND EXPANSION SLOTS

The CPU must talk to planar memory, expansion boards, coprocessor, keyboard, and the like. It communicates with other devices on the motherboard via metal *traces* in the printed circuit. But how can expansion boards be connected to the CPU?

Some computers, like the earlier Macintosh line didn't allow easy expansion. To expand a Mac, the circuit boards must be partially disassembled or modified in ways not intended by the original developers. (Now the new Macs have an expansion bus called NuBus.) Making modifications difficult puts the user at the mercy of the modifier, as virtually all such modifications are done at the expense of the warranty and service agreement, if any. Another disadvantage is that the average Joe/Jane can't do the modifications him/her self. This would be like you having to cut a hole in the wall of your house to find a main power line every time you want to use an appliance. Failing standard interface connectors, you would have to find the power line, then splice the appliance into it to get power for the appliance.

This scenario, as we know, is silly, as we have standard outlet plugs. Any manufacturer who wants to sell me a device requiring electrical power need only ensure that the device takes standard US current, and add a two prong plug. "Upgrading" my house, then (adding the new appliance) is a simple matter: just "plug and play." Many computers adopt a similar approach. Such computers have published a connector standard: any vendor desiring to offer an expansion board for this computer need only follow the connector specifications, and the board will work in the computer.

The first microcomputer, the Altair, had such a standard: it was called the S-100 or Altair bus. The Apple II used a different bus. The PC uses yet another bus, with 62 lines.

The 62 lines mentioned above are offered to the outside world through a standard connector, as mentioned previously. These connectors are called *expansion slots*, as expansion boards must plug into these slots. There are five on the PC and eight on the XT and AT.

There are actually *four* different expansion slots commonly seen in the PC world.

The most common is the simple 62 line, "8 bit" slots found on the original PC and XT. Data transfers can only occur in 8 bit chunks on this bus. Most AT and 386-type machines have at least 2 of these. The PC and XT type clones have *only* these kind.

With the advent of the AT, IBM wanted higher speed capabilities. One way to improve speed is to "widen" the bus. The AT uses both 8 bit expansion slots, which, again, are only wide enough to accommodate 8 bits at a time, *and* 16 bit slots which are wide enough to allow 16 bits to pass.

When Compaq introduced the first 80386 based desktop, it wanted to exploit the 32 bit power of the 386 chip. They included a new 32 bit slot, mainly for use with expansion memory. Some 386 clones have adopted this, but it's not a very significant part of the market.

Finally, IBM changed the rules again on 2 April 1987 when they announced the PS/2 line. In order to facilitate faster data transfer within the computer, and to lower noise levels, the new PS/2 Models 50, 60, and 80 (*not* the 30) have a new bus called the Micro Channel Architecture (MCA) bus. It is completely incompatible with the old bus. Expansion boards made for the PC, XT, or AT will not operate in the PS/2 line. Be sure when you buy expansion boards what kind of machine they are destined for.

2.8. SYSTEM CLOCK

The system clock, mentioned before, is the "metronome" for the computer system. It is implemented on the PC on a chip called the 8284A. It will be covered under the topic of accelerator products to speed up your PC or AT. The 8284A is located near the 8088 on the system board.

2.9. NUMERIC COPROCESSOR

Next to the 8088 on a PC or near the 80286 on an AT is an empty socket. For two years, IBM wouldn't say officially for what the socket was intended. But then they announced what everyone already knew: it was for the Intel 8087, a special purpose microprocessor. The 8087 is a microprocessor that is only good for one class of tasks: floating point numeric operations. If you do a lot of floating point calculations (calculations with numbers that have decimal points), and the software that you use is written to take advantage of an 8087, then such a chip is a good investment. (The classic program that can benefit is Lotus.) They cost about \$150 at this writing, and installation is simple — see the chapter "Circuit Boards and Chips."

The AT needs an 80287: they cost about \$200 - \$300. 386-based machines may use either an 80287 or 80387, depending on how they were designed. 80387s can be found for about \$500.

2.10. POWER SUPPLY

U.S. line current is 120 volts alternating current (AC). The PC, like most digital devices, is set up to use direct current (DC) at 5 and 12 volts. The conversion process is done by a *power supply*. The power supply is the silver or black box to the rear and right of the PC. Power supplies are rated by the amount of power that they can handle: 63.5 watt power supplies are used on PCs, '30 watt power supplies on XTs, and 200 watt power supplies on ATs. Your power supply determines in part how long your computer components last. A good power supply costs less than \$200.

The chapter "The Power Supply" discusses selecting, replacing, and installing a power supply. The power supply can't always cope with environmental conditions, so the chapter also discusses add-on products: surge suppressors, spike isolators, and uninterruptible power supplies.

2.11. KEYBOARD

The PC is useless without an input device, and the keyboard is the input device used by most of us. The keyboard is subject to a number of hazards, however, and needs maintenance — and sometimes replacement. The PC's keyboard actually contains a microprocessor of its own, called the Intel 8048. Taking apart a keyboard isn't hard — re-assembling it *is*. Attacking both problems, and discussing alternative and replacement models is the subject of the chapter "The Keyboard."

2.12. DISPLAY AND DISPLAY ADAPTERS

All peripheral devices need something to communicate between them and the computer. Sometimes these are called "interfaces", "ports", or "adapters." The PC, like most computers, uses video cathode ray tube (CRT) technology to display information for user reception.

To allow the computer to communicate with a display monitor, a display adapter must be inserted in one of the PC's expansion slots. Several display adapters are available:

- The IBM Monochrome Display Adapter and Compatibles
- The Hercules Monochrome Graphics Adapter and Compatibles
- The IBM Color/Graphics Adapter and Compatibles
- The IBM Enhanced Graphics Adapter and Compatibles
- The IBM Professional Graphics Adapter and Compatibles
- The IBM Video Graphics Array
- The IBM Multi Color Graphics Array
- The IBM 8514/A Very High Resolution Graphics Adapter
- Non-IBM High Resolution Graphics Adapters

Each adapter may service one or more types of displays. The major families of displays are the following:

- Monochrome TTL monitors
- Composite Video Monitors
- RGB Monitors
- High Resolution RGB Monitors
- Multiscan Monitors

I have no intention of telling you how to open up and attempt to service one of these monitors — it's dangerous and not particularly cost effective. However, some simple maintenance and troubleshooting techniques can be done with the monitor cover in place. I'll discuss them in the chapter on "Displays and Display Adapters".

2.13. FLOPPY AND HARD DISK CONTROLLER AND DISK DRIVES

The floppy disk drive is an essential peripheral. As a peripheral, it requires an interface card. This card is called a floppy disk controller board. These will, in general, not give you many problems. They *do* fail, however, so you need to know how to recognize this and address it.

The much more fertile ground for failure lies in the floppy disks and floppy disk drives themselves. Floppy drives can require speed adjustment, head alignment, and head cleaning. Speed adjustment and head cleaning can be done simply and cheaply. Alignment can require some specialized equipment and is not always cost effective, but will be discussed later.

Hard disks are a godsend if you've been using floppies — they store more data and are faster — but they bring their own host of problems. Hard disk (IBM calls them "fixed disks") controllers can develop problems, or may contain the key to better hard disk performance. There are a lot of options in hard disk controllers these days, and we will examine those options.

Some hard disk failures — precipitous drops in speed or loss of data — can be addressed at either the controller or the drive level. Problems like head crashes can be avoided with some simple techniques explained in the chapter "Hard Disks." Even the ultimate disaster — a reformatted hard disk — can be reversed in some cases.

The XT and AT use controllers with radically different designs. In general, they're not interchangeable.

2.14. PRINTERS AND PLOTTERS

Despite the electronic age and the paperless office, we still don't believe that it's true until we see it on paper. Printers and plotters produce this "hard copy" output for us. As should be familiar by now, such devices require an interface, generally a *Centronics parallel port* or an *RS232C serial port*. In the case of the parallel port, the interface usually poses few or no problems: serial ports can be troublesome sometimes, but inexpensive diagnostic tools described in "Modems and Communications Ports" can quickly isolate problems.

The actual printer is the greater source of failures: printers employ a large number of moving parts. Some dot matrix and laser printers are very reliable, while some daisy wheel printers can offer no end of problems. The chapter "Printers and Plotters" examines printer maintenance.

Printer ports have names to DOS. DOS calls the first printer port on a PC "LPT1," or "Line PrinTer 1." DOS also supports LPT2 and LPT3.

2.15. SYSTEM CLOCK/CALENDAR

The system clock/calendar keeps the date and time even when the unit is turned off. It is found either on the system board (the AT) or an expansion board (PC or XT family). The vast majority of the time, the only problem from the clock will be with the battery. Clocks are discussed in "Circuit Boards and Chips."

2.16. MODEMS AND COMMUNICATION PORTS

Modems allow computers to communicate with other computers via phone lines. The interface used for most modems is the RS232C interface. It is also called a *serial port* or a *comm port*.

RS232 is a source of many cable problems: either the wrong cable is configured, or environmental problems (electronic noise) cause communication errors. Data communications troubleshooting is an entire class in itself — several classes, in fact — but we cover the essentials in "Modems and Communications Ports."

The DOS names for the communications ports are COM1 and COM2. DOS 3.3 will also allow COM3 and COM4.

3. AVOIDING SERVICE: PREVENTATIVE MAINTENANCE

3.1 INTRODUCTION

The most effective way to cut down your repair bills is by good preventative maintenance. Some of this is common sense. However, many misconceptions exist in the PC community about what makes and doesn't make good maintenance sense. A few factors endanger your PC's health:

- Dust
- Magnetism
- Electromagnetic Noise
- Water
- Temperature

3.2. STEP ONE TO MAINTENANCE: HOW DO I GET INSIDE THE PC?

Some regular maintenance requires disassembly of the PC. We need to make a small digression before getting on to maintenance.

3.2.1. PC Disassembly Steps

1. First, you need a lot of room. Get a cup to store screws and small pieces of hardware.
2. Remove the monitor from atop the PC and put it aside. (Actually, you're only supposed to put monitors on top of the PC if they are 17 pounds or less. The monochrome monitor passes, but most color and RGB monitors do not.)
3. Turn the PC and associated peripherals off.
4. At this point, you're supposed to unplug the power cords, but I never do. Please ensure that the machine's power is turned off.
5. On the back of the PC, there are five (two on pre-1983 PC's) screws. Remove them and put them in the cup. Don't knock the cup on the floor yet: wait for more, smaller screws and hardware. Slide the cover forward and set it aside. If you've never done it before, take this opportunity to put your name on the inside of the case. A small extra bit of security.
6. If you are going to remove the circuit board, detach all connectors to it.

Now you're ready for open-PC surgery.

3.3. AVOIDING DUST

Dust is everywhere. It consists of tiny sand particles, fossil skeletons of creatures which lived millions of years ago, dead skin, paper perfore, and tiny crustaceans called dust mites which live off the other pieces. Dust is responsible for several evils.

First, it sticks to the circuit boards inside your computer. As dust builds up, the entire board can become coated with a fine insulating sheath. That would be fine if the dust were insulating your house, but thermal insulation is a definitely bad thing for computers. You seek, as we will see later, to minimize impediments to thermal radiation from your computer components.

To combat this, remove dust from inside the computer and from circuit boards periodically. A good period between cleaning is a year in a house and six months in an office. A simpler approach is to use the "while I'm at it" algorithm — when you need to disassemble the machine for some other reason, clean the insides while you're at it. A tool which can assist you is the "Mini-Vac", made by Mini-Vac, Inc., at:

Mini-Vac, Inc.
P.O. Box 3981
Glendale, CA 91201
(818) 244-6777

The Mini-Vac is a *nice* product in some applications, but not a *great* product. Compressed air is better in many cases. Just as effective for the case and bracket assemblies is a dust-free cloth wetted with a little water and ammonia (just a few drops). Don't use the cloth on circuit boards — get a can of compressed air and blow the dust off. When you blow dust off boards, be aware of where it is going: if you can have the vacuum cleaner nearby, or take the board to another area, then you'll have better luck.

The second dust evil is that dust can clog spaces, like:

- the air intake area to your power supply or hard disk
- the space between the floppy disk drive head and the disk

To combat the floppy drive problem, some manufacturers offer a floppy dustcover which you put in place when the machine is off. The sad part of this is that you really need the cover when the machine is on. The reason: CRT displays have an unintended, unpleasant, unavoidable (the three U's) side effect: they attract dust. Turn your screen on, and all of the dust in the area drops everything (what would dust particles drop, I wonder?) and heads straight for the display. Some of the particles get side-tracked and end up in the floppy drives. Some vendors say that the way to cut down on dust in floppy drives is to close the drive doors. This is wrong for two reasons: first, the door *isn't* dust-tight. **Second, double sided drives should be stored with the doors open.** This is because the heads do not touch each other, and cannot damage each other, when the drive doors are open. Obviously, this is a moot point for half height drives, which do not allow you to close the doors unless a floppy is in the drive.

A place which creates and collects paper dust is, of course, the printer. Printers should be vacuumed or blown out periodically, AWAY from the computer (remember dust goes somewhere when blown away).

By the way, another fertile source of dust is ash particles. Most of us don't burn things indoors, *unless* we are smokers. If you smoke, fine: just don't do it near the computer. An OSHA study estimated that smoke at a computer workstation cuts its life by 40%. That's \$1200 on a \$3000 workstation.

3.4. MAGNETISM

Magnets — both the permanent and electromagnetic type — can cause permanent loss of data on hard or floppy disks. The most common magnetism found in the office environment is produced by electric motors and electromagnets. A commonly overlooked electromagnet is the one in phones which ring (rather than chirp or beep). The clapper is forced against the bell in the phone by powering an electromagnet. If you absent-mindedly put such a phone on top of a stack of floppy disks, and the phone rings, you will have unrecoverable data errors on at least the top one.

Another source of magnetism is, believe it or not, a CRT. I have seen disk drives refuse to function because they were situated inches from a CRT. X-ray machines in airports similarly produce some magnetism, although there is some controversy here. Some folks say, "don't run floppies through the X-ray — walk them through." Others say the X-ray is okay, but the metal detector zaps floppies. Some people claim to have been burned at both. Personally, I walk through an average of three to four metal detectors per week carrying 3-1/2 inch floppy disks, and have never (knock wood) had a problem.

What about preventative maintenance? For starters, get a beeping or chirping phone to minimize the chance of erasing data inadvertently. Another large source of magnetism is the motor in the printer — generally, it is not shielded (the motors on the drives *are* shielded, for obvious reasons.)

A sad story: a large government agency's data center bought a hand-held magnetic bulk floppy eraser. (I'm not sure why — they weren't a Secret shop, and thus did not have the need.) The PC expert in the shop tested it on a few junk floppies, then turned it off and didn't think about it. The next day, he remembered that he had left it on top of a plastic floppy file drawer. This meant that the eraser, even though turned off, was about an inch from the top of the floppies. He spent the next day testing each of the floppies, one by one. Most were dead. They got rid of the bulk eraser. I'm not sure what they did with the PC expert.

3.5. ELECTROMAGNETIC NOISE

Stray electromagnetic radiation can cause problems for your PC. "Noise" in this context is, of course, unwanted radiation. Strictly speaking, unwanted electromagnetism can be conducted as well as radiated. The types of electromagnetic noise are *electromagnetic interference* and *radio frequency interference*.

3.5.1. Electromagnetic Interference

Electromagnetic interference (EMI) is low frequency radiation. Sources are:

- static electricity (high voltage, low current)
- power surges (temporary current increase)
- voltage spike (temporary voltage increase)
- radiated EMI from motors
- crosstalk between wires and cables

EMI can cause temporary or permanent damage to your computer. Power line problems and their solutions are discussed in the later chapter, "Power Supplies". Static electricity, however, is not a power line problem: it is power generated by you.

3.5.2. Power-Up Power Surges

I'd like to discuss one power-related item here: user-induced power surges. *What* user-induced power surges, you say? Simple: every time you turn on an electrical device you get a power surge through it. **Some of the greatest stresses that electrical devices receive is when turned on or turned off.** When do light bulbs burn out? Think about it — they generally burn out when you first turn them on or off.

The answer? Leave your PCs on 24 hours/day, seven days/week. We've done it at MM&Co. for years. Turn the monitor screen intensity down, or use one of those annoying automatic screen blankers so the monitor doesn't get an image burned into it. There are other advantages, also: see the later subsection "Thermal Shock."

3.5.3. Static

Static electricity is annoyingly familiar to anyone who has lived through a winter indoors. The air is very dry (relative humidity 20% in my house, for example), and is thus an excellent insulator. You build up a static charge, *and keep it*. In the summer, when relative humidities can be close to 100% (I live in a suburb of Washington, DC, a city built over a swamp), you build up static charges also, but they leak away quickly due to the humidity of the air.

You know how static electricity is built up. But here's some details that you might not have known. Scuffing across a shag rug in February can build up 50,000 volts. This is an electron "debt" which must be paid. The next metal item (metal gives up electrons easily) pays the debt with an electric shock. If it's 50,000 volts, why don't you electrocute when you touch the metal? Simple — fortunately, the amperage is tiny. (Power = voltage times current.) Different materials generate *more or less* static. Many people think that certain materials are static-prone, while others are not. As it turns out, materials have what Litton Systems, Inc., calls a triboelectric value. Two materials rubbed together will generate static in direct proportion to how far apart their triboelectric values are. Some common materials, in order of their triboelectric values, are:

- Air
- Human skin
- Asbestos
- Rabbit Fur
- Glass
- Human hair
- Nylon
- Wool
- Fur
- Lead
- Silk
- Aluminum
- Paper
- Cotton
- Steel
- Wood
- Hard rubber
- Nickel and Copper
- Brass and Silver
- Gold and Platinum
- Acetate and Rayon
- Polyester
- Polyurethane
- Polyvinyl Chloride
- Silicon
- Teflon¹

Once an item is charged, the voltage potential between it and another object is proportional to the distance between it and the other item on the table. For instance, suppose I charge a glass rod with a cotton cloth. The glass will attract things below it on the scale, like paper, but will attract more strongly things below paper.

¹Source: Brenner, page 143.

Why does static damage PC components? The chips which largely comprise circuit boards are devices that can be damaged by high voltage, even if at low current. CMOS chips can be destroyed by as little as 250 volts. Even if static doesn't destroy a chip, it can shorten its life. Static is, then, something to be avoided if possible. The easiest way that I use is to ground myself to something metal *which is not the computer's case*. A metal desk or table leg is good². For your firm, however, you may want something a trifle more automatic. The options are:

- raise the humidity with a humidifier
- install static free carpet
- put antistatic "touch me" mats under the PCs

From the point of view of comfort, I recommend the first option strongly. Your employees don't feel dried-out, and the static problem disappears. Raise humidity to just 40% and the problem will go away.

Notice that I did not include some others:

- use antistatic spray
- use an antistatic mat under the PC

I haven't had much luck with either of those, and don't recommend them.

3.5.4. Crosstalk

When two wires are physically close to each other, they can transmit interference between themselves. We're not talking about short circuits here: the insulation can be completely intact. The problem is that the interfering wire contains electronic pulses. Electronic pulses produce magnetic fields as a side effect. The wire being interfered with is touched or crossed by the magnetic fields. Magnetic fields crossing or touching a wire produce electronic pulses as a side effect. (Nature is, unfortunately, amazingly symmetrical at times like this.) The electronic pulses created in the second wire are faint copies of the pulses, i.e. the signal, from the first wire. This interferes with the signal that we're trying to send on the second wire.

Crosstalk is not really a problem when applied to power lines, although I have heard of cases where the alternating current in power lines creates a hum on a communications line through crosstalk. The larger worry is when bundles of wires are stored in close quarters, and the wires are data cables.

There are three solutions to crosstalk:

- move the wires farther apart (not always feasible)
- use twisted pair (varying the number of twists reduces crosstalk)
- use shielded cable (the shield reduces crosstalk)

²A number of articles and books say, "touch the case of the power supply," presumably because it is grounded. I say don't do this, not because you could get a shock, but because *not all "grounds" (the middle pole on the power plug) are grounded*. The resulting surge could find its way back into the PC.

3.5.5 Radio Frequency Interference

Radio Frequency interference (RFI) is high (10 KHz+) frequency radiation. It's a bad thing. Sources are:

- high speed digital circuits, like the ones in your computer
- nearby radio sources
- cordless telephones, keyboards
- power-line intercoms
- motors

Worse yet, your PC can be a *source* of RFI. If this happens, the FCC police come to your place of business and take your PC away. (Well, not really. But they *will* fine you. More on this later.)

RFI is bad because it can interfere with high speed digital circuits. Your computer is composed of digital circuits. RFI can seem sinister, because it seems to come and go mysteriously. Like all noise, it is an unwanted signal. How would we go about receiving a *wanted* RF signal? Simple — construct an antenna. Suppose we want to receive a signal of a given frequency? We design an antenna of a particular length. (Basically, the best length is one quarter of the wavelength. A 30 meter wavelength is best picked up by a 7.5 meter antenna. It's not important that you know that. To know more about it, pick up an amateur radio book.) Now suppose there is some kind of RFI floating around. We're safe as long as we can't receive it. But suppose the computer is connected to the printer with a cable that, through bad luck, happens to be the correct length to receive that RFI? The result: printer gremlins. Fortunately, the answer is simple: shorten or lengthen the cable.

Electric motors are common RFI-producing culprits. I recently saw a workstation in Washington where the operator had put an electric fan (to cool *herself*, not the workstation) on top of the workstation. When the fan was on, it warped the top of the CRT's image slightly. Electric can openers, hair dryers, electric razors, electric pencil sharpeners, and printers are candidates. Sometimes it's hard to determine whether the device is messing up the PC simply by feeding back noise onto the power line (the answer there is to put the devices on separate power lines), or whether it is troubling the PC with RFI.

Your PC also *emits* RFI which can impair the functioning of other PCs, televisions, and various sensitive pieces of equipment. By law, a desktop computer cannot be sold unless it meets "class B" specifications. The FCC requires that a device 3 meters from the PC must receive no more than the following RFI:

Frequency	Maximum Field Strength (microvolts/meter)
30 - 88 Mhz	100
89 - 216 Mhz	150
217 - 1000 Mhz	200

Source: Brenner (p.142)

RFI became an issue with personal computers when the PC came out because IBM had shielded its PC line, and sought to make life a little tougher on the clonemakers. By pushing the FCC to get tough on PCs, IBM had a bit of a jump on the market. Unfortunately, getting Class B certification isn't that hard, and most clones qualify these days: clonemakers have started saying that their machines are "FCC Class B Certified." This has caused the reverse of IBM's original intent, as the FCC certification seems a mark of legitimacy. In reality, FCC certification doesn't necessarily indicate good design, quality components, or compatibility.

Protecting your PC from the devices around it, and protecting the devices from your PC are done in the same way. If the PC doesn't leak RFI, then it's less likely to pick up any stray RFI in the area. Any holes in the case provide entry/exit points. Use the brackets which come with the machine to plug any unused expansion slots. Ensure that the case fits together snugly and correctly. If the case includes cutouts for interface connectors, find plates to cover the cutouts or simply use metal tape.

A simple AM radio can be used to monitor RFI field strength. A portable Walkman-type radio is ideal, as it has light headphones and a small enough enclosure to allow fairly local signal strength monitoring. A cheap model is best — you don't want sophisticated noise filtering. Tune it to an area of the dial as far as possible from a strong station. Lower frequencies seem to work best. You'll hear the various devices produce noises. The system that I am currently using has an XT motherboard, a composite monitor, an external hard disk, and a two-drive external Bernoulli box. The quietest part of the system is the PC: the hard disk screams and buzzes, the Bernoulli makes low frequency eggbeater-like sounds, and the monitor produces a fairly pure and relatively loud tone.

The PC sounds different, depending on what it is doing. When I type, I hear a machine gun-like sound. When I ask for a text search, the fairly regular search makes a "dee-dee-dee" sound.

3.6. AVOIDING WATER AND LIQUIDS

Water is an easier hazard to detect and avoid. You don't need any sophisticated detection devices. Shielding is unnecessary — you just keep the computer away from water.

Water and liquids are introduced into a computer system in one of several ways:

- Operator spills
- Leaks
- Flooding

Spills generally threaten the keyboard. One remedy — the one recommended by every article and book I've ever read on maintenance — is to forbid liquids near the computer. In some shops, this is unrealistic. Another answer is a formed plastic cover, such as:

SafeSkin (\$29.95)
Merritt Computer Products, Inc.
4561 S. Westmoreland
Dallas, TX 75237
(214) 339-0753

It is a plastic skin which fits over the keyboard, protecting it from spills. They offer versions for the various odd keyboards in the PC world.

Flooding sometimes occurs. Don't assume that flooded components are destroyed components. Disassemble the computer and clean the boards (see the chapter "Circuit Boards and Chips"), and be sure to clean the contacts and edge connectors: a pencil eraser cleans contacts well. Blow out crevices with compressed air.

Avoid floods by thinking ahead. Don't store any electrical devices directly on the floor: they'll be damaged when the floor is cleaned. Generally, flooding indoors is under six inches. Be aware of flooding from improper roofing: when installing PCs, don't put one in just under the suspicious stain on the ceiling ("oh, *that* — it was fixed two years ago. No problem now").

3.6.1. Corrosion

Liquids (and gases) can accelerate corrosion of PCs and PC components. Corrosive agents include:

- salt sweat in skin oils
- water
- airborne sulfuric acid, salt spray, carbonic acid

Your fear here is not that the PC will fall away to rust: the largest problem that corrosion causes is oxidation of circuit contacts. When a device's connector becomes oxidized, it doesn't conduct as well, and so the device does not function, or — worse — malfunctions sporadically. Salt in sweat can do this, so be careful when handling circuit boards: don't touch edge connectors unless you have to. This is why some firms advertise that they use gold edge connectors: gold is resistant to corrosion.

You don't believe that you have detectable traces of finger oils? Try this simple experiment. Pour a glass of soda or beer into a very clean glass —preferably a plastic cup that has never been used before. There will be a noticeable "head" on the drink. Now put your finger into the center of the head, just for a second. The head will rapidly dissolve, as the oils damage the surface tension required to support the head. It's the quickest way to eliminate a large head so you can pour a larger glass of beer.

Carbonated liquids include carbonic acid, and coffee and tea contain tannic acids. The sugar in soda is eaten by bacteria who leave behind conductive excrement — like hiring some germs to put new traces on your circuit board. Generally, try to be very careful with drinks around computers.

Don't forget cleaning fluids. Be careful with that window cleaner that you're using to keep the display clean. If your AT is on a pedestal on the floor, and the floor is mopped each day, some of the mopping liquid gets into the AT. Cleaning fluids are very corrosive.

3.7. HEAT

Electronic components have a temperature range that they are built to work within. The PC, for instance, is built to work in the range 60 degrees to 90 degrees (Fahrenheit, of course). Your printers and hard disks may require more tight or more tolerant ranges. Heat burns up components, and a well-designed device — circuit board, motor, or monitor — can transport heat away from itself at the same rate at which it produces it.

3.7.1. Duty Cycles

Not every device is that good, however. Devices are said to have a *duty cycle*. This number — expressed as a percentage — is the proportion of the time that a device can work without burning up. For example, a powerful motor may have a 50% duty cycle. This means that it should be active only 50% of the time. A starter motor on a car, for example, must produce a tremendous amount of power. Powerful motors are expensive to produce, so cars use instead a motor that can produce a lot of power for a very short time. If you crank the engine on your car for several minutes at a stretch, then you will likely damage or destroy the car's starter motor. Floppy disk drive motors are a similar example: run a floppy motor continuously and you'll likely burnout the motor. *Hard* disk motors, on the other hand, run continuously and must be designed with a 100% duty cycle. Disk stepper motors for PC hard disks are not designed to run continuously. Thus, it's okay to leave your hard disk turned on 24 hours/day, but probably not okay to make it read/write data constantly.

Duty cycle is used to describe active vs. nonactive time for many kinds of devices, although it is (strictly speaking) not correct. Some desktop laser printers, for example, will not run well if required to run continuously.

3.7.2. Thermal Shock

PCs are supposed to be operated in the temperature range 60 to 90 degrees F. This is not the range that the PC actually operates in, however. Inside the PC, temperatures may rise to 120 or 130 degrees. The inside of the PC that I use hovers about 10 degrees above the temperature of the room.

This leads to another problem, called *thermal shock*. Thermal shock comes from subjecting components to rapid and large changes in temperature. It can disable your computer due to expansion/contraction damage. The most common scenario for thermal shock occurs when the PC is turned on Monday morning after a winter's weekend. Most commercial buildings turn the temperature down to 55 degrees over the weekend: your office may contain some of that residual chill early Monday morning. Inside the PC, though, it may be still 55. Then you turn the machine on. Inside 30 minutes the PC is warmed up to 120 degrees. This rapid, 65 degree rise in temperature over a half hour brings on thermal shock.

This is another argument for leaving the PC on 24 hours/day, 7 days/week. The temperature inside the PC will be better modulated. To monitor the PC's inside temperature, the Edmund Scientific Corporation sells a digital thermal probe which measures temperature both inside and outside of the PC for \$39.95.

Dual Function Digital Lab Thermometer
Catalog #E36,987
The Edmund Scientific Corporation
101 E. Gloucester Pike
Barrington, NJ 08007
(609) 573-6250

This applies *double* to portable PCs. If your laptop has been sitting in the trunk on a cold February day, be sure to give it some time to warm up before trying to use it.

3.7.3. Other Heat Effects

"Ambient" heat — that is, the heat in the environment around us — affects the PC, as it is harder to shed excess heat into the surrounding environment when the environment has a high heat content (for similar reasons, heat pumps serve better as a means to heat houses above 30 degrees than below 30 degrees). A PC will live longer in an environment that stays in the 70 degree range than one in a non-air conditioned room which might see a few 92 degree days each summer.

Heat aids the corrosion process. Corrosion is a chemical process, and chemical processes roughly *double* in speed when the temperature of the process is raised by 10 degrees Centigrade (about 18 degrees Fahrenheit). Chips slowly deteriorate, the hotter the faster.

Power supplies (see the chapter on them) can be bought which will lower the temperature in your PC by 35 degrees. They cost about \$150. 35 degrees means the chemical reactions occur four times slower.

4. TROUBLESHOOTING PROBLEMS: WHAT TO DO WHEN SOMETHING GOES WRONG

Okay, you vacuum out your PC fortnightly. You clean and adjust your disk drives semi-annually. You have a robot that shoots anyone coming within 50 feet of your PC with food or drink. But, one day, WordPerfect refuses to print your purple prose. How do you proceed?

4.1. GENERAL TROUBLESHOOTING RULES

These rules have kept me out of trouble for a long time. I know they'll be of use to you.

The Rules of Troubleshooting Almost Anything

- (1) I WILL WIN**
- (2) DRAW A PICTURE**
- (3) IDENTIFY COMPONENTS**
- (4) TRUST NO ONE**
- (5) VERIFY WORKING EQUIPMENT**
- (6) REBOOT AND TRY AGAIN**
- (7) REMOVE MEMORY RESIDENTS**
- (8) USE AVAILABLE TEST EQUIPMENT**
- (9) NEVER GUESS**
- (10) OBSERVE LIKE SHERLOCK HOLMES**
- (11) WISH FOR LUCK**

4.2. TROUBLESHOOTING STEPS

All troubleshooting problems can be decomposed into steps. You *must* be methodical, or you will thrash helplessly about and get frustrated. Once you are frustrated, you are *lost*. Once it happens, you can't think straight and you start creating *new* problems.

Following is the method that *I* use. It looks a lot like methods suggested by other people, but it's not the only method. Whatever method you like, stick to it. It's the "this will only take five minutes" repairs that get me in trouble.

Before opening the computer up, do the following:

1. Check the nut behind the keyboard.
2. Check that everything is plugged in: power, monitor, phone lines, printer, modem, and so on.
3. Check the software.
4. Ask: "what am I doing differently?"
5. Check external signs. Make notes of them.
6. Run the diagnostics diskette.

4.3. CHECK FOR OPERATOR ERROR

Operator error is responsible for 93.3 percent of PC failures³. There are lots of things that an operator can do wrong. For example, when I teach hands-on dBase III+ classes, I'll regularly see a student looking bewildered at the screen. "It didn't work," they say. I say, "Did you type in exactly what I told you to?" "Yes," they say. I look at their screen. They *did*, indeed, type in exactly what I told them to this time. Three commands back, however, they miskeyed something and ignored the resulting error message. These are intelligent people: they're just under some stress (they have to pay attention to me *and* the computer), and they miss a detail. It's easy to do. You're an intelligent person. But you're under some stress if you are trying to fix a computer: you probably need the computer for some very important task. Trust me (rule 4 notwithstanding): if you try to repair while panicking, it'll take longer. Take the time. Do it right.

Even worse, sometimes users will (horrors!) prevaricate slightly. "I didn't do anything. It just stopped working." Please note: I'm not one of those techie types whose motto is, "assume that the user is lying." But sometimes it happens.

Another source of operator error is with inexperienced operators. The PC isn't exactly the simplest thing in the world to master. The author of a book titled Computer Wimp: 100 Things I Wish I'd Known Before I Bought My First Personal Computer observes in that book that learning to use a computer system may be the most difficult learning endeavor that a person may undertake in his/her post-school life. (Things like raising kids are undoubtedly tougher, but they're different kinds of learning experiences.) It doesn't take a genius to recognize that the IBM manuals aren't the easiest things to comprehend.

4.4. IS EVERYTHING PLUGGED IN?

I know this sounds stupid, but we've all done it. A friend bought a Hayes Smartmodem 2400 and couldn't get it to work. It accepted commands all right, but could not dial out or receive calls. The phone line was tested with a regular phone and worked fine. He was quite puzzled until he realized that he plugged the phone line into the "out" jack in the modem (intended to be connected to a phone so that a phone line can be shared between the modem and a phone), rather than the "in" jack. Another time, I ran the IBM diagnostic disk on my PC and kept getting "bad address mark" errors on my disk drives. I was all set to spend a lot of money for new drives until I realized that the scratch diskette that I was using was not formatted. (The drives were fine.)

³That's a made-up statistic. But it got your attention and probably isn't far from the truth.

Not only is everything plugged in, but is everything in *tight*? Multiple pin connectors slowly bend under gravity unless the mounting screws are tightened. As someone stretches his/her legs under the desk, a loose power cord could be moved enough to disconnect it, or disconnect and reconnect it. Connectors on the floor take a lot of abuse.

4.5. CHECK THE SOFTWARE

Curse the person who invented memory resident software.

Software problems arrive in several guises:

- operator error (see above)
- keyboard/screen/disk conflicts with memory resident software
- software that doesn't clean up after itself
- software that requires hardware that isn't connected or activated
- buggy software

4.5.1. Resource Conflicts Among Software Packages

Memory resident software, like Sidekick, Metro, Prokey, Superkey, and the like are great. You don't have to exit dBase or WordPerfect to do a simple calculation, make a note, or enter an appointment in an electronic calendar. Just press a particular key combination and Sidekick pops up. You do your work, then press Escape and Sidekick goes back to sleep.

At least, that's the way it's *supposed* to work. Often it doesn't, however. Sometimes you press the key combination and the PC goes to sleep.

When odd things occur software-wise, memory resident applications are the first suspect. Reboot without the memory residents and try the operation again. If the problem goes away, try to reproduce the software failure and contact the manufacturer (see the chapter "The Most Common Problem: Software").

Sometimes you can rearrange the "hot" keys in memory resident applications. Sometimes re-arranging the order in which you load memory residents solves the problem. When you find a workaround, make a note of it so that you and your successors can save time later.

These conflicts are a major reason for the new PC operating system, OS/2. OS/2 creates multiple "virtual PCs," each running a separate program. Sidekick, then, thinks that it has its own machine, and doesn't conflict with WordPerfect's ability to get to the keyboard. This presumes that Sidekick is written to "play by the rules" that OS/2 lays down. Virtually no one played by the PC-DOS rules: how well OS/2 rules are followed only time will tell.

4.5.2. Poorly Terminated Software

Most software is fairly careful to restore your PC to its original state when it ends. Some, on the other hand, isn't so careful. Most software effects are fairly innocuous, like leaving the background of your display blue and the foreground yellow. The simple DOS command "MODE BW80" will fix that. Perhaps you have both a color display and a monochrome display, and the software has left you in the color display: "MODE MONO" will fix that. Occasionally, software will disable Ctrl-Alt-Del (Flight Simulator is one example) or set the disk drives to nonstandard parameters. The only option may be the big red switch or a soft reboot. (See the chapter "Advanced Projects" to see how to install a switch to do hardware resets.)

4.5.3. Hardware-Related Software Faults

For a while, you would occasionally see people writing in to computer magazines about a mysterious bug in the original Compaq Portable. For no reason at all, it would sometimes freeze up, and nothing would save it. Eventually the computer world figured out that the problem was sloppy fingers on the part of the operator.

The Compaq, like all PC compatibles and clones, has a `PrtSc` (print screen) key. Pressing shift and `PrtSc` sends an image of the screen to the printer. If a printer is not present, a PC does several retries and finally decides that there is no printer attached. This is called "timing out." The PC doesn't wait very long — just a few seconds. The Compaq, on the other hand, waits quite a long time. Thus, when the stray finger pressed shift-`PrtSc`, the computer would wait and wait and wait for a printer to be attached. Actually, had the user waited long enough, the computer would have come back.

A number of mysteries can be linked to software that doesn't recover well from disabled or nonexistent hardware.

- Trying to print to a nonexistent printer
- Trying to print to an Epson printer when a C. Itoh is attached
- Trying to print to a printer which is off line
- Trying to display graphics data on a monochrome monitor
- Running a program that assumes (generally for copy protection) that the PC is 4.77 Mhz speed, or that the floppy drive has 41 tracks
- Running a program that needs more memory than the PC contains

4.5.4. Faulty Software

Sometimes the problem is just plain buggy software. Even the most popular programs can misbehave when faced with a full disk, insufficient memory, or some situation that the designer didn't anticipate or didn't test.

Try to make the bug reproducible. If it is a suspected bug in a compiler product, trim off as much of the other code as is possible while retaining the bug. Ideally, a program no longer than 10 lines of code should demonstrate the problem. Then report it to the manufacturer, and other users in your company.

4.6. WHAT AM I DOING DIFFERENTLY?

When I used to work in a large programming shop, junior people would often come to me for diagnostic advice on their code.

"This code worked before, but it doesn't now". "What did you change?" I would ask. "Nothing," they usually replied. "Then *why* did you run it again?" I wondered.

There *has* to be something different. Otherwise, why are you running the software? Was some code changed? It's the first place to look. Did you add a memory board? Some early PC programs (Infocom games, in particular) would refuse to run if more than 512K memory was in the PC. "Illegal Operation," it would complain, and reboot the PC. You'd think that *adding* memory couldn't cause problems — particularly if you'd already checked the memory and found it error-free. In the case of the Infocom games, "Illegal Operation" isn't much of a hint about what's wrong. I barked up a lot of trees until I went back and examined the single change to the system: a memory upgrade.

If you have such a problem, and must convince your system that it has less memory than it actually has, the program disk has a program called `MEMSET.COM`. `MEMSET` will tell your system that it has whatever memory you tell it.

Hardware upgrades can conflict. Suppose, for example, you install an Intel AboveBoard paged memory board and a 3Com LAN board in the same PC⁴. Suppose the AboveBoard uses the same memory window as the LAN board. Most of the time, the AboveBoard is inactive: it only is called to work when a spreadsheet too large to fit into 640K is being worked on. The LAN board, on the other hand, is used periodically to store and retrieve data, and to receive electronic mail. The user boots up alright, and sees that the LAN works fine. She then starts working on Lotus. As long as the spreadsheet is small, no problem. *Then* she begins using the AboveBoard memory. Still no problem. *Then* disaster strikes: someone sends her an electronic mail message. The LAN and the AboveBoard are active at the same time, and the system either freezes up, or the LAN goes south, or the spreadsheet is minced, or some combination of the above.

Software upgrades can bring problems. Perhaps an application running under DOS 2.1 has *just* enough memory space to run: the extra twenty-odd kilobytes required by DOS 3.2 makes it unable to run. A spreadsheet macro that required three minutes to run under version 2 of the spreadsheet might take two hours or not run at all under version 3.

4.7. CHECK EXTERNAL SIGNS

If the computer has indicator lights, what do they indicate? Are all of the lights glowing on the modem? Does the printer indicate "ready?" Is the hard disk squealing or grinding? Does the monitor image look bent? Your drives and other peripherals produce hums, whirrs, and clicks. After a while, these noises become familiar, and any variation in them signals a problem. Pay attention to these signs.

The first step in successful troubleshooting is in isolating the problem component. These signs can point the way.

It is very important to note any signs here. Note what lights are on and off, the positions of switches, etc.

4.8. RUN DIAGNOSTIC PROGRAMS

The IBM PC and AT come with diagnostic programs which can help pinpoint the problem (assuming, of course, that the computer is well enough to run them in the first place). Other computers do not. Various public domain diagnostic programs exist. If you feel like spending money, SuperSoft makes a diagnostic product:

SuperSoft Service Diagnostics
\$225 for 8088/8086 based PCs
\$225 for 80286 based PCs
P.O. Box 1628
Champaign, IL 61820

This runs on any clone or compatible: machine specific versions are available. Alternatively, IBM offers their Advanced Diagnostics for \$175.

Truthfully, the programs are not so valuable to *locate* errors. They mainly make you feel confident that there is nothing wrong with the PC.

⁴This example is not based on actual experienced problems. The boards were arbitrarily picked for the purpose of example.

4.8.1. Running The IBM Diagnostic Programs

The IBM diagnostic disk is a bootable disk: put in drive A:, then reboot the system. It uses a simpler version of COMMAND.COM, so don't put it directly on your hard disk. This text will not dwell on running the diagnostic disk, as it is pretty simple — all menu driven. The IBM manual lacks some information on interpreting error codes, however, so a table of diagnostic error codes are presented in the section "IBM Diagnostic Error Codes."

4.8.2. The Power On Self Test (POST)

A short diagnostic routine runs on the PC every time the PC is powered up. It does a check to see that the basic important hardware exists, and a memory test. It is called the POST — Power-On Self Test, or, in the BIOS listing, POD (Power-On Diagnostics). The documentation given with the PC (at least the documentation given with *my* PC) does not explain the error codes. Hopefully, you will never see them. If you do, however, you will want to know what they mean. They are shown in the section "IBM Diagnostic Error Codes."

The following description is taken from my reading of the assembly language listing for the PC-1 BIOS. It varies slightly for the AT and later PCs.

1. POST tests the 8088: registers and flags. If it fails, the system just halts.
2. POST computes and checks a checksum for the ROM. If it fails, the system just halts.
3. POST checks the DMA (Direct Memory Access) Controller, the 8237. If it fails, the system just halts.
4. Test the 8259 interrupt controller. If it fails, do a long beep and a short beep, then halt.
5. Test the 8253 timer: is it running at the right speed. If it is not, do a long beep and a short beep. The code actually jumps to the same point as for the error message for the 8259, above. This is unfortunate, as it would have been nice if IBM could have added 8 more bytes of code so that the 8253 could have a different error message from the 8259.
6. Perform a checksum test on the BASIC ROMs. If the ROMs fail, long and short beep and halt. Again, this is a bit short-sighted. We can live without BASIC. Locking up the computer seems extreme.
7. Test video. If the 6845 video controller is not present (which either means no video card or a defective 6845), two long beeps and a short.
8. Ask installed adapters if they must initialize themselves. Common examples are hard disk controllers, EGA, and LAN adapters. If the boards respond, "yes," let them initialize. When they are done, return to BIOS startup.
9. Test the CRT interface lines. If horizontal and vertical sweep do not appear, two longs and a short beep. If everything is okay, blink the cursor on the PC.
10. Memory test. Test all system RAM.
11. Keyboard test. Stuck keys detected here. Any keyboard problems are indicated by a "301" code on the screen followed by a short beep. The system is *not* halted.

12. Cassette interface test. If problems, show code "131" and a short beep. Do not halt system.
13. Test diskette adapter and drive A:. Attempt to reset the drive and activate the drive motor. If problems, show code "601" and a short beep. Go to Cassette BASIC.
14. Determine how many printers, serial ports, game ports, etc. are attached. Issue a short beep, load the boot record, and transfer control to it.

4.8.3. IBM Diagnostic Error Codes

Here are the meanings of the error codes which occur in the POST and the IBM Diagnostics:

POST Audio Messages

No beep, nothing happens	Power supply or not plugged in
Continuous beep	Power supply bad
Repeating short beep	System board
1 long beep, 1 short beep	System board
1 long, 2 short beeps	Failure or lack of display adapter/cable
1 short beep, blank screen	Failure or lack of display adapter/cable
1 short beep, no boot	Floppy Drive Adapter Failure

IBM Hardware Diagnostics/POST Messages

CODE	DESCRIPTION
01x	Undetermined Problem Errors
02x	Power Supply Errors
1xx	System Board Errors:
101	Interrupt Failure
102	Timer Failure
103	Timer Interrupt Failure
104	Protected Mode Failure
105	Last 8042 Command Not Accepted
106	Converting Logic Test
107	Hot NMI Test
108	Timer bus test
109	Direct Memory Access Test Error
121	Unexpected Hardware Interrupts Occurred
131	Cassette wrap test failed
161	AT Battery failure
162	AT Setup Info incorrect (re-run SETUP)
163	Time and Date not set (run SETUP)
164	Memory Size Error (run SETUP)
199	User Indicated Configuration Not Correct
2xx	Main Memory (RAM) Errors
201	Memory Test Failed
202	Memory Address Error
203	Memory Address Error
(see below to decode this error message)	

-
- 3xx Keyboard Errors**
- 301 Keyboard Error. If followed by a number, the number is the scan code of the key in question.
 - 302 User indicated error from keyboard test or AT keylock locked
 - 303 Keyboard or system unit error
 - 304 CMOS does not match system
- 4xx Monochrome Monitor Errors**
- 401 Adapter memory, horizontal sync frequency test, or video test failed
 - 408 User indicated display attributes failure
 - 416 User indicated character set failure
 - 424 User indicated 80 x 25 mode failure
 - 432 Parallel port test failed
- 5xx Color Monitor Errors**
- 501 Color Adapter Memory, horizontal sync frequency test, or video test failed
 - 508 User indicated display attributes failure
 - 516 User indicated character set failure
 - 524 User indicated 80 x 25 mode failure
 - 532 User indicated 40 x 25 mode failure
 - 540 User indicated 320 x 200 mode failure
 - 548 User indicated 640 x 200 mode failure
- 6xx Diskette Drive/Controller Failures**
- 601 Adapter or Drive failed POST
 - 602 Diskette test failed: boot record is not valid
 - 606 Diskette verify function failed
 - 607 Write-protected diskette
 - 608 Bad command diskette status returned
 - 610 Diskette initialize failed
 - 611 Timeout
 - 612 Bad NEC chip on diskette controller
 - 613 Adapter failed DMA test
 - 621 Bad seek
 - 622 Bad CRC found
 - 623 Record not found
 - 624 Bad Address mark
 - 625 Bad NEC seek
 - 626 Diskette data compare error
- 7xx 8087 or 80287 Math Coprocessor Errors**
- 9xx Printer Adapter Errors**
- 1101 Asynchronous (RS232) Adapter Failure (COM1)
 - 1201 Asynchronous (RS232) Adapter Failure (COM2)
- 13xx Game Port Failure**
- 1301 Adapter test failed
 - 1302 Joystick test failed
- 14xx Printer Errors**
- 1401 Printer Test Failed
 - 1402 Dot Matrix Printer Test Failed

15xx SDLC Adapter (Mainframe Connection) Failures

1510	Failure of 8255 port B
1511	Failure of 8255 port A
1512	Failure of 8255 port C
1513	8253 timer 1 did not reach terminal count
1514	8253 timer 1 stuck on
1515	8253 timer 0 did not reach terminal count
1516	8253 timer 0 stuck on
1517	8253 timer 2 did not reach terminal count
1518	8253 timer 2 stuck on
1519	8273 port B error
1520	8273 port A error
1521	8273 command/read timeout
1522	Interrupt level 4 failure
1523	Ring Indicate stuck on
1524	Receive clock stuck on
1525	Transmit clock stuck on
1526	Test indicate stuck on
1527	Ring indicate not on
1528	Receive clock not on
1529	Transmit clock not on
1531	Data Set Ready not on
1530	Test Indicate not on
1532	Carrier Detect not on
1533	Clear To Send not on
1534	Data Set Read stuck on
1536	Clear to Send stuck on
1537	Level 3 Interrupt Failure
1538	Receive Interrupt results error
1539	Wrap data miscompare
1540	DMA channel 1 error
1541	Error in 8273 error check or status reporting
1547	Stray interrupt level 4
1548	Stray interrupt level 3
1549	Interrupt presentation sequence timeout

16xx Terminal Emulation Errors (32xx, 5520, 525x)**17xx Hard Disk/Disk Controller Errors**

1701	POST error
1702	Adapter failure
1703	Drive failure
1704	Drive or Adapter failure: cannot be determined
1780	Drive 0 failure (drive C:)
1781	Drive 1 failure (drive D:)
1782	Adapter failure
1790	Drive 0 failure
1791	Drive 1 failure

18xx Expansion Chassis Failures

1801	POST error code
1810	Extender Card Failure
1811	Extender Card Failure
1816	Extender Card Failure
1820	Receiver Card Failure

1821	Receiver Card Failure
1812	Address or Wait State Failure
1813	Address or Wait State Failure
1819	Wait Request Switch Set Incorrectly

19xx 3270 PC Communications Controller Failures

20xx BSC Adapter (Mainframe Connection) Failures

2010	Failure of 8255 port B
2011	Failure of 8255 port A
2012	Failure of 8255 port C
2013	8253 timer 1 did not reach terminal count
2014	8253 timer 1 stuck on
2015	8253 timer 0 did not reach terminal count
2016	8253 timer 0 stuck on
2017	8253 timer 2 did not reach terminal count
2018	8253 timer 2 stuck on
2019	8273 port B error
2020	8273 port A error
2021	8273 command/read timeout
2022	Interrupt level 4 failure
2023	Ring Indicate stuck on
2024	Receive clock stuck on
2025	Transmit clock stuck on
2026	Test indicate stuck on
2027	Ring indicate not on
2028	Receive clock not on
2029	Transmit clock not on
2031	Data Set Ready not on
2030	Test Indicate not on
2032	Carrier Detect not on
2033	Clear To Send not on
2034	Data Set Read stuck on
2036	Clear to Send stuck on
2037	Level 3 Interrupt Failure
2038	Receive Interrupt results error
2039	Wrap data miscompare
2040	DMA channel 1 error
2041	Error in 8273 error check or status reporting
2047	Stray interrupt level 4
2048	Stray interrupt level 3
2049	Interrupt presentation sequence timeout

21xx Alternate BSC Adapter Failures (same as above, but with "21" prefix rather than "20")

2201 PC Wiring Cluster Adapter Failure

2401 EGA Failure

2901 Color Dot Matrix Failures

3301 Compact Printer Failures

5. CIRCUIT BOARDS AND CHIPS

If the problem isn't software, if everything's plugged in, and it isn't something obvious like a burnt-out motor, a broken wire, gummed-up printer, or an imploded display, then it's probably a circuit board. Step one is to identify the faulty part. Diagnostic software and diagnostic hardware procedures are identified in this section. Step two is repair. There are two levels of repair here: board level and chip level.

Sometimes, software can diagnose which component is faulty. In some cases, however, hardware diagnostic techniques are required. The following section discusses that.

Board level repair involves swapping bad boards. This is simple and cost effective. You just replace the bad board with a good one, then throw away the bad board. The next level involves repairing the bad boards, which is done generally by removing and replacing bad chips, capacitors, transistors, and the like. (As far as I know, no one repairs chips, capacitors, etc.)

As important and troublesome as repairing board-type problems is installing new boards. Personally, this creates more trouble for me than do problems with existing boards.

5.1. BURN-IN

Most solid state components (i.e., electronic things without moving parts) tend to be very reliable if they live past the first few days. "If it works, it'll work forever." Because of this, initial testing of circuit boards is valuable. Such testing is called *burn-in*.

When a new component like a memory board is installed, it should be tested for a 24 to 72 hour period. Some diagnostic program — generally a simple programming loop of some kind — retests the component over and over again. Some components have warranties of only 90 days, so the burn-in increases the probability that a bad component will be spotted in the warranty period.

Some manufacturers claim to do burn-in for you. You should *still* do it yourself. Why? They burned it in in the factory. Since then it has been subjected to the rigors of shipping.

5.2. BOARD IDENTIFICATION: BOARDS FOUND IN THE PC

Generally, troubleshooting computers consists of identifying and replacing a damaged circuit board. The first step to this is knowing what boards exist in the PC.

The following table summarizes the boards found in a PC, and their approximate replacement costs:

XT/PC System Board	\$105
Floppy Disk Controller Board	31
Hard Disk Controller Board	150
Hard/Floppy Controller (AT)	160
Display Controller Board:	
Color/Graphics Adapter (CGA)	74
Monochrome Display Adapter (MDA)	79
Monochrome Graphics Adapter	74
Enhanced Graphics Adapter (EGA)	179
Professional Graphics Adapter (PGA)	\$\$\$
Memory Expansion Board	49

Asynchronous Communications Adapter (serial port)	39
Printer Adapter (parallel port)	49
Real Time Clock/Calendar	49
Multifunction Board (clock, printer, serial, memory)	85
More "exotic" boards:	
3278 Emulator Board	500 - 1200
Local Area Network (LAN) Board	200 - 1000
Expanded (Paged) Memory Board	125

5.2.1. Maintenance Considerations For The PS/2 Family

Currently, the selection of available boards for the PS/2 line is considerably smaller than that for the PC family. There are, however, few boards necessary. The display controller, disk controller, parallel and serial ports, and clock/calendar are built right onto the motherboard (oops—*planar* board).

Boards will still be required. LANs, extra memory, micro/mainframe, data acquisition and the like will require expansion boards, and they will become available slowly.

In some senses, this is a good thing. Less boards to worry about. Less DIP switches to fuss with. But it also has a very serious implication for troubleshooting. Replacements for the new IBM motherboards are quite expensive — over \$1400 for the Model 50. **And you can't fix the boards — you must buy replacements.** Even if you could locate the faulty chip on the board, it is likely a custom IBM chip ("gate"). IBM won't sell you a replacement gate. It doesn't matter how good or bad an engineer you are, fixing a faulty PS/2 motherboard is just plain impossible.

A PS/2 is then a risky proposition: it's not likely to fail, but if it does, it'll cost lots of money.

Don't think you can't maintain them — everything in this text applies to PS/2s as well as PCs. But if the planar board fails, you can't do anything short of buying a new one, at least until someone builds a cheap clone of the planar board.

5.3. INSTALLING NEW CIRCUIT BOARDS

As I said above, you will probably have to configure and install new circuit boards more often than you will have to fix existing boards. This section looks at configuration and installation.

5.3.1. Circuit Board Removal And Insertion

The physical process of installing or removing circuit boards is straightforward.

Removing an expansion board:

1. The power must always be off.
2. The back of the board has a *mounting bracket*. The mounting bracket is attached to the back of the PC case with a screw. Remove the screw. Put the screw in the cup with other screws (remember the cup? It's all right to knock it over now.)

3. If the mounting bracket has an interface connector on it, like a serial interface or a parallel interface, and something is connected to it, disconnect it.
4. Grasp the board with both hands, one on the top front edge and the other on the top rear edge.
5. Rock the board back and forth, pulling up gently.
6. The board will come out.
7. When you remove a board, you should fill the empty space on the PC box's rear bracket with a bracket insert. This keeps RFI to a minimum. If inserting a board, save the bracket insert.

The insertion process is just the reverse. If you are inserting a new board, ensure that you configure it first, if necessary (see below).

Generally, it makes *no difference* which slot boards are inserted in. The only criterion for board placement is ease of cable routing. There are two exceptions. (1) don't use the slot on the XT that is closest to the power supply. It requires slightly different boards. This does not apply to clones. (2) If installing an accelerator board install it as close to the CPU as is possible short of using the bad XT slot.

If the board that you inserted/removed was a memory board, don't forget to inform the system via the DIP switches that system memory has changed.

5.3.2. Installing And Configuring Boards: DMA, IRQ, and I/O Ports

Sometimes the problem stems from a conflict between one board and another board in the computer: two boards trying to use the same resource. This might happen if you tried to install two RS232C interface boards, both set for device name COM1:. Ordinarily, this problem will show up when the second board is installed, but not always. It is, then, appropriate at this point to discuss board installation.

To add a board, it must first be configured, then installed, and finally tested. We have already discussed installation in the beginning of this chapter. Test with a burn-in diagnostic program. The next few sections discuss configuration.

Configuration involves setting jumpers or DIP switches to select exactly what services the expansion board will provide. Typical configuration requirements are:

- tell an expansion memory board how much memory is on the board
- tell a serial port whether it is COM1: or COM2:
- tell a printer port whether it is LPT1, LPT2, or LPT3
- select DMA channels on a board (more on this later)
- select IRQ lines on a board (more later)
- select I/O ports on a board (ditto)
- indicate whether a floppy is drive A: or B:

PS/2 owners, you will be spared DIP switches and jumpers. PS/2 expansion boards are intended to do all that with software. You will be able to configure a board without touching it.

In many cases, the board is pre-configured at the factory to the correct settings, but not always. It is hard for the manufacturer to know what the proper settings should be for three items: I/O ports, DMA channels, and IRQ lines.

5.3.3. I/O Ports

How is information actually transferred from a peripheral, like a serial port, to the microprocessor? Well, the microprocessor can read and write memory directly, so maybe we could fool the 8088 into thinking that the serial port is actually a location in memory. Then we could send information out the serial port by reading that "memory address," and we could read information by storing data to that "memory address."

Such an approach is called *memory mapped I/O*. It is done in the Motorola line of products, like the 68000 (found in the Macintosh), but not in the Intel line of products. Intel uses I/O ports. The way these work is that for every memory address in the low 64K there is a "shadow" I/O address. Thus, there are about one million memory locations on the PC, and about 65,536 "I/O locations" on the PC's processor, the 8088. Unfortunately, IBM only implemented the bottom 1024 ports. Devices designed to use port locations above 1023, or 03FF hexadecimal, will not work on the PC. For economy's sake, the same wires used to address memory are used to address I/O ports. The question of *which* task it serves at the moment is answered by a separate wire. Also, the I/O port approach requires a separate command from the MOV ("move") assembler command used to read/write memory on the 8088. I/O ports are accessed by IN and OUT commands in assembler. On a Motorola chip, reading from and writing to an I/O port would be done with the same command as would be used to read/write memory.

Each device that must communicate with the 8088 has an I/O port or ports. This port cannot be used by any other device. This means that, for example, if you have a motion sensor talking to your PC via port 200, and the PC talks to the printer interface via port 200, then *neither* the motion sensor nor the printer will function properly. This is where the jumpers and DIP switches come in. Realizing that they cannot know what ports are currently in use on your computer, manufacturers give you a choice of possible ports. Moving the jumpers or DIP switches allow the choice. (Note: when the I/O ports are changed, the software usually must be notified of these changes.)

We said that the PC only implements 1024 ports. The bottom 256 (hex locations 000 to 0FF) are only available to components on the system board. Plug-in expansion boards must use the 768 top locations (hex 100 to 3FF). Therefore, expansion boards will only allow you (if they are designed properly) to set your I/O addresses somewhere between 100 hex and 3FF hex.

If troubleshooting is your job, then it is probably a good idea to keep a roster of PCs and busy I/O ports. The table on page 34 tells the common ports taken up on the PC.

5.3.4. DMA Channels And IRQ Lines

Transferring data from a device into the computer via the CPU can be very slow, so some devices have the power to write data into the computer memory directly, without CPU intervention. This is called DMA, Direct Memory Access. The PC has a single DMA controller chip, the 8237. It allows up to four DMA channels. One — channel 0 — is required for dynamic memory refresh.⁵ The floppy disk controller usually employs channel 2. The hard disk controller uses channel 1. 3 is unused in general.

⁵Ever wonder what "dynamic memory refresh" is? (If not, return to the text.) There are two kinds of memory: dynamic and static. Dynamic sounds better than static, but it isn't. When you tell a static RAM something, it remembers it until you turn off the power or change it. Dynamic RAMs, on the other hand, forget whatever you tell them within 4 milliseconds. The PC is designed to drop everything and do a RAM refresh every 3.86 ms. This takes 5 clocks out of every 72, or about 7% of the PC's time. Of course, if the CPU is doing a lot of INs, OUTs internal calculations, or the like, then you don't notice the slowdown, as the whole idea of DMA is to work in parallel with the CPU. Wouldn't static RAMs make a slightly faster computer? Yes, but they're more costly. Compaq uses them in the Deskpro 386.

Multiple DMA channels can run in parallel. The AT has two DMA controllers, and thus eight DMA channels to the PC's four. Note that not all AT clones have this second DMA controller. This will no doubt yield problems later, as the market for 80286 products matures.

Common I/O Port Uses in The PC Family

Hex Address Range	User
00 - 0F	DMA Controller 8237
100 - 1F	DMA Controller
20 - 21	Interrupt Controller 8259A
22 - 3F	Interrupt Controller
40 - 43	Timer 8253
44 - 5F	Timer
60 - 63	8255 Peripheral Controller (8088 only)
60 - 6F	Keyboard
70 - 7F	NMI (non-maskable interrupt) mask register
80 - 83	DMA Page Registers
0A	NMI mask register
C0 - DF	DMA Controller 2
F0 - FF	Math Coprocessor
1F0 - 1F8	Hark Disk Controller
200 - 20F	Joystick Controller
210 - 217	Expansion Chassis
2F8 - 2FF	COM2 Serial Port
300 - 31F	Prototype Card
320 - 32F	Hard Disk Controller
330 - 337	Bernoulli Box Controller
378 - 37F	LPT1 Printer Port
380 - 38F	SDLC Card
3A0 - 3AF	BSC Card
3B0 - 3BF	Monochrome Adapter
3D0 - 3DF	Color/Graphics Adapter
3F0 - 3F7	Floppy Disk Controller
3F8 - 3FF	COM1 Serial Port

Note: **Bold-faced** entries apply to the AT.

Common DMA Channel Uses in the PC Family

Channel	Use
0	Dynamic RAM Refresh
1	Hard Disk Controller
2	Floppy Controller
3	Unused

To get the CPU's attention, Interrupt ReQuest (IRQ) lines are used. The PC's bus implements lines 2 through 7. They are prioritized, line 2 more important than line 7. When a line is activated, the processor drops everything and loads a special subroutine written to service the particular interrupt.

Common IRQ Line Uses in The PC Family

Device	Interrupt Line
Memory	0
Keyboard	1
Unused	2 (used by some EGA's)
COM2	3
COM1	4
Hard Disk	5
FDC	6
(6 is floppy and hard disk on AT)	
LPT1	7

As interrupt 2 is unused, many "odd" devices like LAN boards make use of it.

A caution. Some boards don't have jumpers and DIP switches. This means that *there is no way to get them to work with conflicting boards*. For example, a client that I regularly visited had installed an IBM 5251 (System 36 terminal emulator) board and an old Quadram Quadboard in a PC. The printer port on the Quadboard and the terminal emulator wanted the same resource — which one, I'm not sure. In any case, neither had jumpers. One board had to be thrown away. Moral: find out if the expansion boards that you buy have adjustable DMA, IRQ, and I/O ports.

I hesitate to mention this, but sometimes device conflicts can be solved by doing surgery on the boards. Just lobotomize the chips that are performing the function which you wish to defeat. An example which I have seen a couple of times is in serial ports. A client wanted me to set up a multifunction board in a PC with clock, memory, printer and serial ports. He already had a board installed which provided both serial ports COM1: and COM2:. The jumpers on the multifunction board allowed me to set the multifunction board's serial port to either COM1: or COM2:, but not disable it altogether. What to do? A chip called the 8250 UART (Universal Asynchronous Receiver/Transmitter) is the heart of most serial ports. I found the 8250 on the multifunction board and removed it. The problem was eliminated. *Please don't try this unless you understand what you are doing.*

5.3.5. ROM Addresses

Added to I/O ports, DMA channels, and IRQ lines is a fourth source of conflict: ROM addresses. Some controller cards (like EGAs and hard disk controllers) require some ROM on-board to hold some low level code. For example, many hard disk controller cards contain code to do a low-level format of the hard disk: just load DEBUG and type G=C800:5. (Please don't even think of doing this before you read the chapter on hard disks unless you know what you're doing.) The XT controller board's ROM starts at C800:0000. As before, a possibility exists that two different boards may require some software on-board, and if the two boards *both* try to locate their ROM at the same location in the PC's memory address space, neither one will work. Fortunately, some boards include jumpers to allow you to move the start address of the ROM.

5.3.6. A Hint For Connecting Cables And Edge Connectors

Expansion boards in the PC are often controller boards for external devices like disk drives or displays. Cables connect the boards to the displays. A common short cable type is a ribbon cable.

Most ribbon connectors are symmetrical with respect to one or more axes. They could plug in one or two ways. Plugging in a connector upside-down will usually damage the device, or the controller board, or both. Most cables are "keyed" — that is, a connector is modified so that it cannot be plugged in incorrectly. Some, however, aren't, and so the following bit of information is valuable.

Ribbon cables are usually blue, white, or gray. One wire is a different color — often red. This wire connects to pin 1 of the connector. This information has saved me more times than I care to recount. I learned it after incinerating a hard disk. (If you're going to make mistakes, you might as well learn from them.)

Edge connectors can become corroded. Use a pencil eraser to clean them (don't hold the board over the PC while doing this).

Occasionally, rough handling can scratch and remove a circuit board trace. This can be repaired by soldering a short (*As short as possible!*) wire across the cut in the trace. If unsure (traces can be faint or thin to begin with), use your ohmmeter. Set to low ohms (Rx1 on the dial), then put the probes on either side of the suspected cut. The meter should read "0" if everything's okay.

5.3.7. Power Requirements

Sure, you've heard about how PC disk drives draw a lot of power, imperiling your power supply and PC, but what about those expansion boards — any worries there? Surely not, many people say: no moving parts. Actually, memory boards draw a considerable amount of power. According to a Mostek datasheet, a 256K chip draws 0.412 watts maximum. A bank of them (nine, recall) is 3.7 watts. 64K chips draw 0.33 watts, so a bank of nine draw 3.0 watts. A full complement of 640K would require $3.7 + 3.7 + 3.0 + 3.0$ equals 13.4 watts at peak — 21% of a 63.5 watt PC power supply! By comparison, some hard disks' peak power consumption is 1.2 watts. This gets worse if you have an older PC or XT that only accepts 64K chips. 10 banks of 64K chips would draw 30 watts peak.

It may be difficult to get power requirement figures for some expansion boards, so use the following rough test: run the PC for a while with the board installed. Then touch the board. A hot board draws a lot of heat. **If your expansion boards are running hot, it's time for an upgraded power supply!** (See the chapter "Power Supplies.")

5.3.8. A Configuration Example

To underscore these concepts, let's look at a simple installation. Suppose you want to install on your system a Bernoulli Box, a high density, high speed floppy disk which offers the same performance as a hard disk. This involves putting a Bernoulli controller in the PC. The controller documentation mentions DIP switches which allow setting the I/O port, but nothing about DMA or IRQ.

Transferring large blocks of data without DMA would be very slow (it even slows floppies), so the board *must* use DMA. A call to Iomega (the manufacturer) yields the information that the board uses DMA channel 3, and no IRQ lines. The lack of IRQ lines makes some sense: why would the Bernoulli interrupt

the processor?⁶ Iomega also tells us that the DMA channel cannot be changed. Fortunately, they have chosen a generally unused DMA channel.

The manual outlines several I/O port possibilities. The factory switch settings use 330-337 hex, which does not conflict with anything. Install the board, and it works the first time. As I have said before, a log of this information makes configuration easier. It's no good knowing that a new board needs IRQ 7 if you don't know if some currently-installed board uses it.

5.3.9. PS/2 Configuration Notes

Those of you with PS/2 Model 50, 60 or 80 may never have to deal with as much preceding pain and suffering.

The Personal System line includes the new Micro Channel Architecture (MCA) bus, an improvement in some ways over the old PC bus. First, no DIP switches. PS/2 machines have 16 IRQ lines — plenty for any application.

You may know that the AT uses battery-backed memory (called "nonvolatile" memory) to retain some configuration information. This does not assist in installing new expansion boards. In the PS/2s, however, this memory remembers IRQ, DMA, and I/O port information. When the machine is powered up, it checks each board against the nonvolatile configuration information. If a new unknown board is detected, the user will be requested to run a setup program.

There's been a lot of hype in the popular press about how you'll never have to worry about configuration again with the PS/2. It is not really true: you still have to worry about two boards not using the same IRQ lines, but fixing the problem is easier — you run a setup program rather than pull the board out and set DIP switches. The extra DMA and IRQ channels will make the configuration job easier, also.

By the way, IBM wasn't the first manufacturer in the PC business to do this. Orchid Technologies has been doing it with their expansion boards since the beginning of 1986.

5.4. PROBLEMS WITH EXISTING BOARDS: COMMON FAILURE POINTS

Most component problems boil down to either environmental trouble, damage due to mishandling, or faulty manufacture. The most common ailments are the following:

- socketed chips creep out of their sockets due to expansion and contraction (rock the chips)
- bad solder joints can disconnect or cause short circuits (resolder)
- weak components fail under heat
- PC board traces can be scratched or lift off
- dirt and dust (see earlier) build up heat
- edge connectors or chip pins corrode
- RFI/EMI impairment

⁶Actually, a disk drive might need to interrupt a processor if it were part of a multitasking system. If task 1 requested a sector, then the data request could be forwarded to the disk controller. While the controller was waiting for the disk drive to return the information, the main processor could do work on task 2 until the controller signalled (via an interrupt) that the data was ready. DOS isn't written like that, however.

5.5. FINDING AND REPLACING BAD BOARDS

You've run the diagnostics, and they tell you that an internal board is defective. Now on to repair. In most cases, the recommended procedure is to replace the board — generally, you don't repair boards, but sometimes you can.

5.5.1. Identifying The Problem Board: Symptoms and Possibilities

There are three basic approaches to troubleshooting when the machine's down. Here I present them.

The first approach: assume that there is only one problem. Ideally, you have two machines, one sick and one well. A simple strategy here is to swap boards, one by one. Each time you swap a board, turn on *both machines* and note which machines are currently well. Ideally, you would like to induce the problem from the originally sick PC to the originally well PC.

Another approach is to strip the PC down to the bare essentials — motherboard and power supply only — and turn it on. You will hear a "beep" if everything is okay. Then add a board, perhaps the display. You should get a blinking cursor, and of course "301" and "601" errors (remember? see last chapter.)

Step	Expected Outcome
Motherboard, Pwr Supply	1 long, 2 short beeps
Add display board	blinking cursor, 301, 601, 1 beep
Add keyboard	blinking cursor, 601, 1 beep
Add disk controller	blinking cursor, 601, 1 beep
Add floppy A:	blinking cursor, 1 beep, boot
Add floppy B:	same, then try B:
Add hard disk control	boot from hard disk
... any other boards ...	

Try it sometime, on a functioning PC. Note the outcomes. Do it *before* you must work with an ailing PC.

The third approach: in many cases the troublesome board can be identified without even running the diagnostics. As discussed above, a board can malfunction due to a hardware problem within it, or it can malfunction due to conflicts with another board. If the problem allows the PC to successfully boot up, the problem may have an obvious cause. If not, consult the following.

Use the following symptoms to point you to the bad board.

Symptom:

On power up, no response. No blinking cursor on the screen. The A: drive light does not come on.

Possibilities:

1. Is the PC getting power? Is it plugged in? Is the fan whirring?
2. Check the intensity/brightness setting on the display.
3. Did you install a new board? Some boards have some computer instructions in ROM (the EGA and hard disk controllers are two examples) which grab control of the computer *before POST even occurs*. If something messes up in them, the PC never even gets a chance. Try the computer without the new board. Did it boot up okay? If so, round up the usual suspects: I/O ports, DMA channels, IRQ lines, and ROM start addresses (see later in this chapter). If *they* are okay, note the date on the ROM chip on the new board and call the manufacturer. You may have a ROM with an old (translation: *buggy*) version of the controller code.

4. It may be the system board. Test the system board test points (see later section by that name), check that the power supply connector to the system board is in place, and ensure that nothing is shorting the system board.

Symptom:

One long beep and two short beeps, drive A: light goes on, no display.

Possibilities:

Most likely display board. Type "dir b:" and ENTER at the keyboard. If the drive B: light comes on, the system is responding — you just can't see it. Swap the display board. If you don't trust yourself to type "dir b:" blind without errors, insert the program disk into drive A: then reboot. The PC will make a distinctive sound.

Symptom:

Blinking cursor, fan whirring, drive A: light comes on and stays on.

OR

Blinking cursor, fan whirring, "601" appears on the screen, drive A: light comes on and stays on or does not come on at all. Cassette BASIC may appear.

Possibilities:

1. Check the cable between the floppy controller and the floppy drive.
2. Check that the floppy is terminated properly (see the chapter on floppy drives).
3. Check the power test points on the floppy drive.
4. Swap the controller board.
5. Swap the A: and B: drives, and move the terminator from the A: drive to the B: drive (see the chapter on floppy drives).
6. Boot disk is damaged.
7. Insufficient memory to run your application.

Symptom:

System boots fine. Printer will not print.

Possibilities:

See the section on printers.

Symptom:

Screen clears, "PARITY CHECK 1" or "PARITY CHECK 2" appears.

Possibilities:

See the section on memory. OR... too little power. Did you just install a power sucker like a hard disk? It may show up in memory errors.

Symptom:

Keyboard does not respond at power-up, or "301" appears.

Possibilities:

1. The system board may be at fault. Test keyboard test points (see chapter on keyboards).
2. Is the keyboard plugged in?
3. Is it an AT keyboard on a PC? Vice versa?
4. Are you leaning on the space bar accidentally?
5. Is this a turbo AT? Some ATs are so fast that they get confused during a reboot (ctrl-alt-del) unless you take your fingers off the keyboard very quickly.

Symptom:

System cannot boot off hard disk.

Possibilities:

1. Is this a new hard disk? See "preparing a hard disk" in the chapter on hard disks.
2. First, see if it can boot *at all*. Try to boot off a floppy in drive A:. If it will, will it then read C:? If so, perhaps you have developed a bad sector in the middle of your boot record, hidden files, or COMMAND.COM. See the section on hard disks to handle this. If this is the first sign of trouble from the hard disk, BACK UP THE ENTIRE DISK TODAY. Before you do anything else. Then see the chapter on hard disks.
3. Check cables to controller. If the hard disk is external, check its power cord and see if its power switch is enabled.
4. Is the disk making unusual noises? See the chapter on hard disks.

5.6. FINDING AND REPLACING BAD CHIPS

Sometimes, you can be a real hero and actually fix a circuit board. Most times you can't. Most maintenance shops do very little board repair, returning the boards to the manufacturer. Board repair can be time consuming and therefore expensive, not a great idea when most boards are under \$100 - \$200. You'll spend a minimum of an hour and a maximum of forever fixing a board.

This is not to say that it should never be attempted. The most prominent chip level installations/troubleshootings are for:

- Memory chips
- Replacement Microprocessors
- Coprocessors
- ROMs (Read-Only Memories)

In general, finding bad chips is done in one of the following ways:

- software identification of chip malfunction
- temperature testing
- digital probe and pulse testing
- use of specialized (and expensive) signature analyzers
- exhaustive chip replacement

Software identification is possible when the chip is not so vital to the system that the system can't run without it. For instance, if the 8088 or 80286 microprocessor goes on vacation, there is no way to run a diagnostic program to begin with. Many memory problems can be tested with software, however, as well as some problems with some major chips, like the 8284 clock chip and the 8237 DMA controller. The ROM chips can be tested with a checksum, if the correct checksum value is known.

Some faulty chips can be traced through their temperature or lack thereof. A properly functioning chip should be slightly warm to the touch — warmer, at least, than after a night of deactivation. A completely cool chip, particularly a large one (large chips tend to run warmer) is probably dead. Similarly, a very hot chip — finger burning variety — is probably dead.

Finally, there is the brute force method. Get a digital probe and pulser and test each chip, one by one. Sure-fire, but slow.

5.6.1. Identifying Chips

A chip can be described by its function, its identification number, and its manufacturer.

Function: a chip may be something as simple as just four NAND gates, each with two inputs (called a quad two-input NAND chip), or as complicated as a microprocessor. Physical size is some indication of a chip's complexity, but is not tremendously important.

Identification Number: when a manufacturer designs a new chip, it gives it an identification number. For example, an 8088 is a particular microprocessor chip, and a 7400 is a quad two-input NAND chip. Generally, chip designs are patented, so another manufacturer must be licensed before it can offer a chip which it did not design. Intel designed the 8088, but Advanced Micro Devices (AMD) probably makes more of them than Intel does, now. I do not know who developed the first 64K x 1 dynamic RAM chip, but, virtually everyone in the chip business makes them now.

Prefixes and suffixes may be added to a chip's ID number. The suffixes refer to the package that it is in (usually "DIP" — Dual Inline Pin, which is indicated by the suffix AN), or the temperature range ("S" is military specification, "N" is normal). Thus, "7400AN" refers to a 7400 chip in a DIP package. A two letter prefix refers to manufacturer. Some manufacturer codes are HD for Hitachi, WD for Western Digital, DM for National Semiconductor, R for Rockwell. Another code may be present, like "B8544." This means (ignore the B) that the chip was made in the 44th week of 1985. Some vendors, like Intel, do not put a date code on their chips. Instead, they put a serial number on the chip: a motherboard in front of me has a chip with serial number L5450275. If you have a 80286 or 80386 based computer, it's probably a good idea to examine and write down your serial number. Periodically, bugs will pop up and certain serial numbers will be recalled. Understand that there is no one single "8088" chip: it has gone through several revisions.

5.6.2. Limitations: How Far Should I Go In Isolating Bad Chips?

First, understand that manufacturers are making it tougher to work with chips. More and more, we are seeing the use of *surface mount* printed circuit boards. These are flatly impossible to work on without special equipment — they are built for robots to work on. Chip replacement is being phased out.

Testing chips is not really economical, or in most cases even feasible.

The average system board has on the order of 50 to 100 chips on it, depending on brand and version. Ideally, we would always like to isolate the bad chip on the circuit board, replace it, and put the board back in service. As we will see, this is often possible. However, in some cases, it just isn't. Chips are composed of what are called *gates*: NOT, NAND, and OR gates are common.

Imagine that you are trying to test some mythical chip that has four gates. To test this chip, we will first check gate 1, then gate 2, gate 3, and 4.

Sounds easy, eh? Not when you realize that each 64K memory chip contains about 256,000 gates. Yes, professional test equipment exists, but it is *very* expensive. When would this make sense? If you have a board — like a serial or printer port — composed of a few relatively simple chips. In the case of the serial port, about eight chips are used: seven simple chips and the 8250 chip. In that case, I'd just replace the 8250, *then* check the other seven chips by hand.

The last alternative is to replace the chips, transistors, and capacitors on the troubled board, one by one. This can, obviously, be expensive and time consuming.

5.6.3. Locating Chips With Heat Problems: Intermittent Failures

As mentioned before, bad chips can sometimes be identified by their temperature or lack thereof.

A worse situation is when the device works *sometimes*, then fails after a while. (Generally just before you are going to save your work.) Heat can make a marginal component stop working, and the marginal component *may not seem unduly warm*. In this situation, how do you locate the bad component? By controlled application of heat and cold.

First, start up the device — PC, modem, whatever — with the cover off. Then blow warm (*I did not say hot* — 130 degrees will do nicely) air onto the circuit board: use a hair dryer. The intermittent chip will fail eventually. Now get a can of component coolant (Radio Shack part number 64-2321, for example) and direct the cold blast directly onto the suspect chip. If there is no suspect chip, start at the big ones. Now try to restart the device. If it starts, there's a good chance that the cooled chip is the bad one. If it doesn't restart, try another chip. Keep notes.

5.6.4. Soldering And Desoldering

A few words about soldering.

It isn't hard, but it *does* take some practice. **If you have never soldered before, the place to learn is not your PC!** Get some practice elsewhere, or just decide not to bother with soldering tasks. If you *do* want to learn, Healthkit has an excellent "Soldering Self-Instruction" kit. It's \$20 and comes with a simple printed circuit project. It comes with a fair amount of information and I like it quite a lot.

We seek to minimize the amount of soldering required. More on sockets later.

The trick with soldering is to heat up *both* components to the desired temperature, then apply the solder. You want just enough heat to melt the solder, but not so much as to destroy whatever it is that you're soldering.

Soldering irons come in various powers or wattages. For PC work, you want a low power iron, like a pencil iron under 50 watts.

- use a pencil iron under 50 watts.
- use a 60/40 solder with rosin core (*not* acid core), 1/32 inch width.
- do not apply the tip more than 10 seconds — this should be more than sufficient.
- for desoldering, *do not* use solder suckers or vacuum bulbs. They can build up static charges. Use wire braid instead.
- Remove the board first. Don't try to solder things on/off boards which are installed in the PC.
- When replacing chips or transistors, socket them first (see next section).
- Buy a solder "jig" (sometimes called a "third hand") so that you have enough hands to the board, the soldering iron, the chip, and the desoldering tool. Edmund Scientific (see earlier) sells one for \$25.95 called the "Extra Hands Work Station."
- If replacing a diode, transistor, or capacitor, draw a picture of how the original is installed. Memory (the human kind) gets faulty when faced with the normal frustration of soldering. It doesn't matter which way you insert a resistor.
- When desoldering a chip from a circuit board, don't desolder each pin in order. This builds up too much heat in one area. Jump around. Use a heat sink. Alternatively, use a solder tip designed for DIP packages.
- Many system boards are now "four layer" boards. They are very tough to work on competently with the usual inexpensive equipment.

5.6.5. Chip Sockets And Chip Insertion/Removal

Chips need not be soldered directly to the motherboard. Chip *sockets* are available ("socket" like light bulb sockets). As long as you're removing and replacing a chip, think about installing a socket for the replacement. Most chips are soldered directly onto the printed circuit board.

Advantages of Socketed Chips:


- easy removal and replacement
- no fear of damaging the chip with solder heat when installing

Advantages of Directly Soldered Chips:

- save money on socket
- no worry about chip creep

As mentioned before, a socketed chip's problem may be no worse than that it has crept far enough out of the socket to impair electrical connection. An early step in troubleshooting is to push all socketed chips gently back into their sockets.

Whether installing a socketed or directly-mounted chip, you must be sure to install it with the correct side up. Chip pins are numbered with the furthest pin on the left hand side labelled 1, then counting down on the left side and finally up the right side (counterclockwise):



A chip can, then, fit into a socket in one of two ways. Install a chip backwards, and you generally destroy the chip. So pay attention when installing.

The top of the chip generally has a notch (indicated above as "n") to orient you when installing. To make things even easier, most circuit boards are designed so that all of the chips face the same way.

5.7. CPU CHIPS

In the Introduction, we discussed CPU chips. In some cases, you will have to remove and replace these chips in order to upgrade or fix a problem. Here is some more technical information about them.

5.7.1. Common PC CPUs

Several CPU chips are used in the PC market. The following table summarizes their characteristics.

Maker	Model	Max Speed (Mhz)	Word Size (bits)	Data Path (bits)	Memory (MB)	Examples of Computers Using
Intel	8088	8	16	8	1	IBM PC, XT, Portable, Jr.
Intel	8086	8	16	16	1	IBM PS/2 Model 30, Compaq Deskpro
Intel	80C86	8	16	16	1	Toshiba 1100+
Intel	80186	8	16	16	1	Tandy 2000, 3Com 3Server
NEC	V20	10	16	8	1	Kaypro 2000
NEC	V30	10	16	16	1	Kaypro PC, NEC Multispeed
Intel	80286	16	16	16	16	IBM PC/AT, PS/2 Model 50 & 60
Intel	80386	24	32	32	2K	Compaq Deskpro 386, PS/2 Model 80

Some of these chips are more or less powerful. Some can actually allow you to improve the throughput of your existing PC. For instance, if you have a PC or PC clone, you can replace your 8088 chip with a V20 chip and see about a 5 -20% increase in performance. The upgrade cost is only about \$20. I'll discuss microprocessors more in the chapter "Circuit Boards and Chips."

5.7.2. The 8088

The 8088 is older and less powerful than the 80286, and is used in the PC. It comes in what is called a "40 pin DIP package," which means a rectangular plastic case with two rows of 20 pins. DIP stands for "Dual In-line Package." Older 8088's are called 8088-1, as they can only run at low speeds (5 Mhz or slower). Newer "turbo" PC/XT clones may run at 6.66, 7.16, or 8.0 Mhz. To do this, they use the 8088-2, which is rated at up to 8 Mhz. The 8088 is the equivalent of about 40,000 transistors.

5.7.3. The 80286

The 80286 is a newer chip, designed by Intel in 1981. Its package is a square of plastic called a PGA, or Pin Grid Array package. The package has an inner and an outer square of pins: there is thus room for 68 pins. It packs a lot more power into a small package than the 8088 does: the 80286 is the equivalent of about 130,000 transistors in about the same volume. Because of this, the 80286 runs hotter and may require extra cooling provisions such as a *heat sink*. Heat sinks are small caps which fit on top of a chip and enable the chip to better dissipate heat generated by the chip.

5.7.4. CPU Speeds

Computers run to the beat of a clock, like a beginning piano player plays to the beat of a metronome. If you set the metronome too fast for the beginning piano player, he/she will become confused and the music won't come out right. Similarly, if you set the clock rate of a CPU too high, it will malfunction. You won't *damage* the chip: the computer just won't function properly. Part of the design of a computer like the PC includes determining a clock rate.

CPU clocks generally "tick" more than a million times per second. A clock which ticks at exactly one million times per second is said to be a *one Megahertz* clock, abbreviated 1 Mhz. The Apple II used a 2 Mhz clock. The original PC and XT use a 4.77 Mhz clock. The AT originally used a 6 Mhz clock, but IBM offers a version which runs at 8 Mhz. Clone makers offer computers with clock speeds up to 16 Mhz. All other things being equal, a faster clock means faster execution and better performance.

5.7.5. Word Size

Any computer can be programmed to manipulate any size number, but the bigger the number, the longer it takes. The largest number that the computer can manipulate in one operation is determined by its *word size*. This is either 8, 16, or 32 bits.

Think of it this way: if I ask you, "what is 5 times 6," you answer, "30," immediately — you did it in one operation. If I ask, "what is 55 times 66," you will do a series of steps to arrive at the answer. 55 is larger than your word size. 5 isn't.

5.7.6. Data Path

No matter how large the computer's word size, the data must be transported into the CPU. This is the width of the computer's "loading door." It can be 8, 16, or 32 bits. Obviously, a wider door will allow more data to be transported in less time than will a more narrow door.

5.7.7. Memory Size

You can't just keep adding memory to your PC indefinitely. A particular chip can only *address* a certain size of memory. For the oldest chips, this amount was 65,536 bytes — a 64K memory. The original PC's CPU can address 1024K, or a megabyte (MB). Other, newer chips can address even more. The 80386 can address gigabytes (billions of bytes).

5.8. TROUBLESHOOTING THE SYSTEM BOARD

If I haven't convinced you not to attack the board level problems yet, let me try three more arguments: first, as I said before, replacement motherboards can be found for about \$120. Second, the motherboard is not difficult to replace, but it is the most time-consuming board to replace. That means that you must go through the roughly 30 minute installation or removal process every time that you try something, like replacing the 8284. That could lose you hours and *still* not answer your problem. Finally (I saved the best for last), for about \$200 you can get a replacement motherboard that makes your PC run *twice as fast*.

5.8.1. Configuration With SETUP And DIP Switches

As the system board is a circuit board, it's a good bet that it requires configuration information, and it does. Check these first if you're having trouble installing a new device onto the board, or installing a new system board. This is contained in the two DIP switch packages on the PC system board, or contained in nonvolatile RAM on the AT system board. If you're not changing anything then the DIP switches probably won't give you any trouble. The nonvolatile RAM *can*, however, give you some trouble if its battery dies.

You tell your PC what equipment it has with two banks of DIP switches on the system board. The DIP switches tell the PC:

- how much memory is installed
- what kind of display is installed
- how many disk drives are installed
- if a numeric coprocessor is installed

5.8.2. PC System Board DIP Switch Settings

Use the tables on page 48 to set the PC's DIP switches. (The AT does not have DIP switches — more on this later.) Use a ballpoint pen or small screwdriver to change the switch settings. If you have set the switches correctly, but the PC refuses to recognize the settings, be aware that **sometimes DIP switches are defective**. (It's happened to me.) To test this, remove the system board and test the switches for continuity with an ohmmeter.

If you don't feel like looking up these tables, your program disk has a program called DIPSET.EXE which will draw the correct DIP switch configuration.

5.8.3. Using the AT SETUP Program

Rather than repeat the DIP switch process with the AT, IBM decided to allow the user to input the information into a nice, menu-driven program called SETUP. SETUP then stores the configuration information into a small bit of RAM. As RAM forgets when power is removed, this small bit of RAM has a battery to keep its power at all times. When the battery runs down, the AT starts acting aphasic (look it up, but it means what you expect it to).

If this is happening to the AT, you need a new battery. New batteries can be bought in the \$17 - 30 range. One such vendor is:

AT Replacement Battery (\$17)
Purner Enterprises
563 San Mario
Solana Beach, CA 92075
(619) 481-1311

The only troublesome thing about the SETUP program is when it wants to know what kind of hard disk you have. On the old XT-type controllers, you set DIP switches and jumpers to select drive type. On the AT controller, the info goes into the battery powered memory. Unfortunately, the drive type number is not particularly descriptive. The table "Built-In AT Drive Types" on page 49, describes each of the 14 types of drives that the AT can accommodate. If you have a questionable controller (such as if you have a clone), run the program HDTABLE.COM on your program disk to see the values that the controller contains.

PC Non-Memory DIP Switch Settings

Coprocessor

Installed SW-1 #2 OFF

Not Installed SW-1 #2 ON

Disk Drives

No Floppies SW-1 #1 ON #7 ON #8 ON

1 drive SW-1 #1 OFF #7 ON #8 ON

2 drives SW-1 #1 OFF #7 OFF #8 ON

3 drives SW-1 #1 OFF #7 ON #8 OFF

4 drives SW-1 #1 OFF #7 OFF #8 OFF

Monitor Type

None SW-1 #5 ON #6 ON

CGA, 40 x 25 Video SW-1 #5 OFF #6 ON

CGA, 80 x 25 Video SW-1 #5 ON #6 OFF

Monochrome Display SW-1 #5 OFF #6 OFF

IBM PC Memory DIP Switch Settings

In the figures below, "1" indicates "ON", and "0" indicates "OFF."

Memory Size (KB)	SW-1		SW-2							
	#3	#4	#1	#2	#3	#4	#5	#6	#7	#8
16	1	1	1	1	1	1	1	1	1	1
32	0	1	1	1	1	1	1	1	1	1
48	1	0	1	1	1	1	1	1	1	1
64	0	0	1	1	1	1	1	1	1	1
96	0	0	0	1	1	1	1	1	1	1
128	0	0	1	0	1	1	1	1	1	1
160	0	0	0	0	1	1	1	1	1	1
192	0	0	1	1	0	1	1	1	1	1
224	0	0	0	1	0	1	1	1	1	1
256	0	0	1	0	0	1	1	1	1	1
288	0	0	0	0	0	1	1	1	1	1
320	0	0	0	1	1	0	1	1	1	1
352	0	0	0	1	1	0	1	1	1	1
384	0	0	1	0	1	0	1	1	1	1
416	0	0	0	0	1	0	1	1	1	1
448	0	0	1	1	0	0	1	1	1	1
480	0	0	0	1	0	0	1	1	1	1
512	0	0	1	0	0	0	1	1	1	1
544	0	0	0	0	0	0	1	1	1	1
576	0	0	1	1	1	1	0	1	1	1
608	0	0	0	1	1	1	0	1	1	1
640	0	0	1	0	1	1	0	1	1	1

BUILT-IN AT DRIVE TYPES

(NOTE: Your AT may not have all 47 types. IBM added them over time.)

DRV #:	Cylinders:	Heads:	Precomp:	Capacity:	Examples:
1	306	4	128	10.6	Seagate ST412, ST212, ST112
2	615	4	300	21.4	Tulin TL226, Qume R200
3	615	6	300	32.1	Tulin TL240, Rodime RO206
4	940	8	512	65.4	Atasi 3080
5	940	6	512	49.0	
6	615	4	no	21.4	Seagate ST4096
7	462	8	256	32.1	Quantum Q540
8	733	5	no	31.9	Seagate ST4038
9	900	15	no	117.5	Maxtor XT-1140
10	820	3	no	21.4	Micropolis 1302, Vertex V130
11	855	5	no	37.2	Vertex V150
12	855	7	no	52.0	Vertex V170
13	306	8	128	21.3	Seagate ST425, Rodime RO204
14	733	7	no	44.6	
16	612	4	no	20.3	
17	977	5	300	40.5	
18	977	7	no	56.7	
19	1024	7	512	59.5	
20	733	5	300	30.4	
21	733	7	300	42.5	
22	733	5	300	30.4	
23	306	4	0	10.1	
25	615	4	0	20.4	
26	1024	4	no	34.0	
27	1024	5	no	42.5	
28	1024	8	no	68.0	
29	512	8	256	34.0	
30	615	2	615	10.2	<--note the skip from 30 to 35. It's not an error, it's what I found on an AT.
35	1024	9	1024	76.5	
36	1024	5	512	42.5	
37	830	10	no	68.8	
38	823	10	256	68.3	
39	615	4	128	20.4	
40	615	8	128	40.8	
41	917	15	no	114.1	
42	1023	15	no	127.3	
43	823	10	512	68.3	
44	820	6	no	40.8	
45	1024	8	no	68.0	
46	925	9	no	69.1	
47	699	7	256	40.6	

5.8.4. Test Points For The System Board

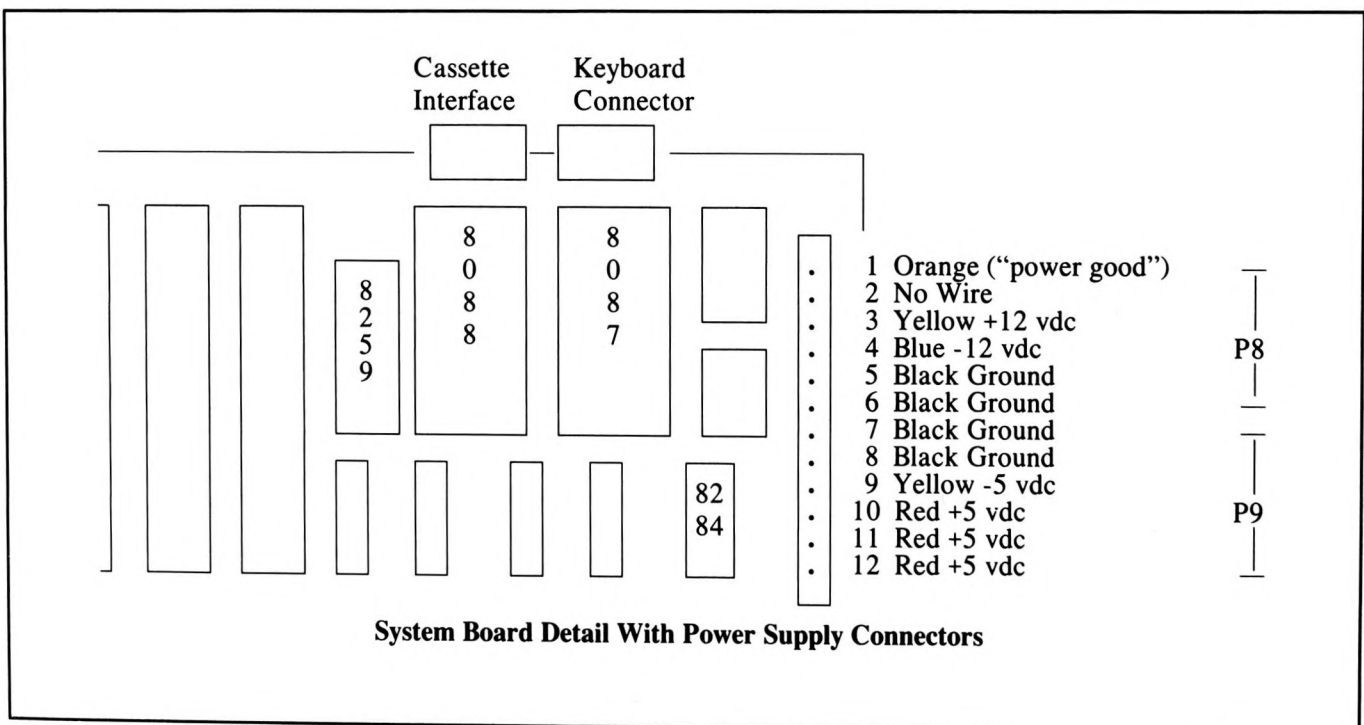
The following is a
section for experts

The remaining part of this section is pretty technical. If you want to skip it, go ahead. As I said before, this whole exercise is probably not cost effective.

If you're still here, let's get at it. We'll assume that you've done all of the basic "is it plugged in?" type tests by now.

First, check that it is receiving power. The motherboard receives power through the power strip in the northeastern corner of the board. (This assumes that you are looking at the board so that the memory chips are closest to you and the expansion connectors, the 8088, and the keyboard connector are farthest from you.) The connector is actually two six pin connectors lined up one atop the other. They are generally labelled P8 and P9.

P8 has only five wires, and connects above P9, which has six wires. See the figure below, "System Board Detail With Power Supply Connectors".



The power supply lines — the yellow, blue, and red wires — should be tested against a ground — any of the black wires. In the process, all of the black wires should be tested. Following are the specifications for these lines:

Wire Color	Rated Voltage (volts)	Acceptable Range (volts)	Current Range (amps)
Yellow	+12	+8.5 - +12.6	0.0 - 2.00
Blue	-12	-8.5 - -12.6	0.0 - 0.25
Red	+5	+2.4 - +5.2	2.3 - 7.00
Yellow	-5	-4.5 - -5.4	0.0 - 0.30

Then disconnect the motherboard from the power supply and do a resistance test. The tests are conducted on the power supply pins. The table "System Board Resistances", below, contains the *minimum* resistance for each connection. If the measured values are less than this, the motherboard is definitely faulty.

System Board Resistances

Com Lead (Black Probe)	Vom Lead (Red Probe)	Minimum Resistance (ohms)
8	10	0.8
8	11	0.8
8	12	0.8
5	3	6.0
6	4	48.0
7	9	17.0

Look inside the chassis. Is there something metal — a screw, iron filing, connector — shorting out two traces? Remove the system board and clean the area under it on the chassis, and blow the system board clean, then reinstall it and the expansion boards.

If it is a newly inserted motherboard, did you use the spacers correctly? Some clone motherboards require plastic spacers to electrically insulate the motherboard from the chassis.

Check traces. Was one cut when you inserted a board? Solder a jumper wire to fix it.

That's about as far as you can go without special devices. If you have access to logic signal measuring devices (like an oscilloscope), try the following approach:

1. Check the power supply inputs (already discussed).
2. Check the clock. Pin 19 on the 8088 should be receiving a clock pulse every 209.5 nanoseconds. If this is a "turbo" PC, the clock pulses should occur according to the following table:

Clock Rate (Mhz)	Time Between Pulses (ns)
4.77	209.5
6.66	150.0
7.16	139.5
8.00	125.0
9.54	104.8
10.00	100.0

If the clock is wrong, and power is right, then try replacing the 8284 (it is socketed, so it's easy). Zero in on the traces around it — any problems? Another possible problem is the crystal (the flat silver can lying nearby on its side). It's soldered in, but there's only two connections and replacements are cheap.

3. Check the reset circuitry. When the PC is first powered up, the RESET line on the 8088 must be activated for four clock periods. Then the 8088 knows where to go to get its instructions. (Do you really want to know? Okay — it's FFFF:0000 hex.) If the reset circuitry — simple though it is — messes up, the 8088 wanders around in circles. So, how is RESET created?

Once again, we look to the 8284 clock chip. The game plan is the following.

The power supply is activated. Once it's humming, it activates its orange wire, "Power Good." "Power Good" is attached to the "reset" line on the 8284. If the "Power Good" line flickers, this "pulls the trigger" on the 8284. That then makes the 8284 generate a 4 clock RESET signal to the PC. So, test points are (1) Power Good (pin 1 on the power supply, connected to pin 11 on the 8284) and (2) RESET (pin 10 on the 8284, which is connected to pin 21 on the 8088.)

If it *still* doesn't work, get it repaired or, better, buy a new one. Don't forget to burn it in in either case.

End of Expert Section

5.9. SEMICONDUCTOR MEMORY

RAM chips are the chips that you are most likely to have to mess with. The most common ones are so-called "dynamic" RAMs. They come commonly (as of this writing) in sizes of 64K and 256K, with the occasional 16K or 1 megabit chip used. In general, dynamic RAMs used in PCs and ATs are 64K x 1 bit or 256K x 1 bit. This means that each chip stores 64K *bits* or 256K *bits*. Thus, to store 256K *bytes*, eight chips are needed — 8 bits equals 1 byte. That is why each 64K or 256K group is a row of chips. Actually, nine chips are used. The extra chip is for error checking with parity. A 256K chip is the same physical size as a 64K chip.

Another important criterion is *access time*. Access time is how fast the chip can respond to a request: if the CPU says, "get me the value in location 31,824", what's the longest that the memory chip will require to respond? Access time is measured in nanoseconds (ns). The slowest memory you'll ever see on the PC is 250 ns. 100 ns and faster are being used in 10 Mhz AT compatibles.

The official names of memory chips reveal two things: their capacity (16K, 64K, 256K, 1024K) and their access time. Chip names look like "41XX-YY", where XX=size and YY=access time in tens of nanoseconds. Some examples:

4164-25	64K,	access time = 250 ns
41256-10	256K,	access time = 100 ns
4164-15	64K,	access time = 150 ns

Recall that the number may be hidden in a longer ID code, like "SN4164-20N".

In case the point has not been made before, use the slowest memory that your machine can use. A regular PC can work with 250 ns or faster memory. You could yank out all of your PC memory and replace it with 100 ns memory — faster memory is perfectly acceptable — but there is no point. The faster memory is more expensive and does nothing. *The computer does not run faster just because you installed faster memory.*

5.9.1. Reading Memory Error Messages

There are three ways to find memory problems: (1) run a rigorous memory test periodically, (2) get a memory error message when the PC does the POST, or (3) get a "PARITY CHECK" error while in the middle of an application. I'll show you how to design your own rigorous memory check later. For now, let's see how to read the cryptic POST messages for memory problems.

The POST, recall, runs a memory test each time that you power up the PC. How does the test know how much memory the system has? You tell it by setting DIP switches on the PC system board, or by running the SETUP program on the AT.

If you've just installed new memory, and you get a PARITY CHECK message, the first thing to check is: did you reset the system board switches correctly? If you installed 512K on a PC and told the system that 640K exists then you'll get a parity error every time. This is not true, by the way, for all clones. Some clones don't offer any complaint if you lie to them about how much memory is installed. XTs really don't pay attention to the DIP switches to determine how much memory they have. If, on power up, the XT finds a problem in the 384K block of memory, it will not issue an error: it will just assume that only 320K memory exists.

Recall that the POST memory error message is a code 201. A POST "201" error will be accompanied by a four digit code, like:

1020 201

followed quickly by the familiar "PARITY CHECK 1" or "PARITY CHECK 2" error. The four digits will help you locate the bad chip. Here's how to do it.

First, "PARITY CHECK 1" implies that the bad memory is on the system board — called "planar memory." "PARITY CHECK 2" implies that the bad memory is on an expansion board.


5.9.2. Reading Parity Check Messages

Parity Check 1 indicates a problem with system board memory. A pretty common cause for this is mis-set DIP switches on the system board: if you tell the DIP switches that the board contains 256K but it actually has only 192K, you'll get an error. The PC will try to use the nonexistent memory and come up with PARITY CHECK 1. So, step 1 is to ensure that the DIP switches are set correctly.

Parity Check 2 just means that the bad memory chip is on an expansion board.

Suppose you have an error message like "3008 201." The first two digits identify the bank that the bad chip is in, the last two digits identify the bad chip within the bank. "201" just tells you that you have a memory error, and that's no surprise. The following section (thru page 57) will show you how to decipher these messages.

HOW TO READ A MEMORY ERROR MESSAGE:



This identifies the actual bank of memory

To understand this, read the following pages "Reading Memory Errors: Deciphering the First Two Digits."

This identifies which chip is bad within that bank. As you look at a bank of chips, you will notice there are nine chips. The second two-digit code identifies which chip is bad within that bank. Reading from left to right, the chip ID's are as follows: 00, 01, 02, 04, 08, 10, 20, 40, 80.

5.9.3. Other Causes of Memory Errors

If a chip seems okay, but causes errors, consider the following rules:

1. Never mix chip speeds in the same row. Do not install a row with some 200 ns chips and some 150 ns chips, even if 200 is the speed required. Remember that if your computer requires 200 ns chips, then a speed of 200 or faster (like 150, 120, 100, or 75 ns) will work fine. You can have one row of 75's, a row of 120's, a row of 150's, and a row of 200's for all it matters — *just don't mix chip speeds in the same row!* (Actually, if your computer only requires 200 ns chips, there's no point in spending the extra money for faster, more expensive chips.)
2. Whenever possible, avoid mixing *manufacturers* in the same row. I flatly cannot explain this, but I have seen cases like the following. I have two rows of chips — one entirely Mitsubishi, another entirely Toshiba — that work fine. Then I mix the rows. Errors occur. Restore the rows, and the problem disappears.
3. Remember to check chip seating: has the chip crept out of its socket?
4. Is there enough power? Insufficient power can cause parity errors. This one can be a real pain, as it waits for some large disk access to trigger the parity error, like when you try to save your data to disk.
5. Improperly shielded sources of RF noise can alter memory, causing parity errors.





READING MEMORY ERRORS: DECIPHERING THE FIRST TWO DIGITS

First, identify which KIND of motherboard you have. There are three.

PC Motherboard Type 1: Original PC 1981 - early 1983

Marking on motherboard: "16 - 64K CPU"

16K banks on the motherboard: Total memory on motherboard = 64K

Memory Chips	Bank Size	1st 2 Digits on "201" error
	16k	00
	16k	04
	16k	08
	16k	OC





Examples: "0420", "0C08" could be on the motherboard
 "1020", "7080", "5001" could not

**READING MEMORY ERRORS: DECIPHERING THE FIRST TWO DIGITS
(CONTINUED)**

PC Motherboard Type 2: PC (1983 - 1987) and XT (1983 - 1986)

Marking on motherboard: "64 - 256K CPU"

64K banks on the motherboard: Total memory on motherboard = 256K

Memory Chips	Bank Size	1st 2 Digits on "201" error
	64k	00
	64k	10
	64k	20
	64k	30





Examples: "1020", "0008" could be on the motherboard
"7080", "5001" could not

READING MEMORY ERRORS: DECIPHERING THE FIRST TWO DIGITS (CONTINUED)

PC Motherboard Type 3: XT (1986 - 1987) and clones (1985 - present)

Marking on motherboard: "256 - 640K CPU"

First two banks are 256K banks, last two are 64K banks. Total = 640K

Memory Chips	Bank Size	1st 2 Digits on "201" error
	256k	00, 10, 20, or 30
	256k	40, 50, 60, 70
	64k	80
	64k	90

Examples: All memory errors are on the motherboard, as all memory is on the motherboard.

5.9.4. Handling Parity Check A Little More Gently: NMI.COM

Granted, we want to address memory problems directly and immediately. But it's no fun losing four hours' work to a stray cosmic ray. Locking up the machine and announcing, "PARITY CHECK," seems a mite drastic. One bad bit out of 4.8 million is survivable. That's why we include NMI.COM on your program disk. It tells the PC to refrain from locking up the system when a parity check is detected, and instead beeps. You can save your work, *then* track down the memory error with lower blood pressure than if you had lost a half day's work.

To use NMI.COM, just include it in your AUTOEXEC.BAT file. Hopefully you will never hear from it.

DO NOT ignore parity beeps! They'll get you eventually.

5.9.5 Tips On Installing Memory Chips

A lot of manuals include some really scary instructions for installing memory chips. I've probably installed a gigabyte or so myself (all right, maybe it's just a tenth of a gigabyte), so here's what works for me.

First, as always, be aware of static electricity. Your acrylic sweater and the soles of your shoes are powerful producers of it. If I'm worried about the static level of an area, *I take my shoes and socks off*. I know it sounds a little bizarre, but I kill very few chips that way – probably less than a half dozen in all the time I've been installing memory.

The problem is that when you pick up memory chips you may have different charges on different fingers of the same hand. Spread out a sheet of aluminum foil and dump the chips so that they sit on their legs on the foil. The foil, recall, conducts electricity. Put all five fingers on the foil. Now the charges are equalized. You can now pick up a chip with impunity.

Second, the memory chips go into sockets in a circuit board. As before, don't insert memory into boards while the power is on. The legs of the chip will spread a bit too much to fit into the socket. Bend the pins so that they are at right angles to the chip body.

Third, insert the chip. Orient it so that its notch is facing in the same direction as the notches of the other chips on the circuit board. Just position the legs over the socket holes (be careful here – don't break or bend the legs) and firmly push the chip into the socket.

If you make a mistake, or must remove old chips, use a small "tweaker" screwdriver. Use it to gently pry up one end of the chip, then work the screwdriver completely under the chip. It will come up easily. Make sure you put the screwdriver under the chip, *not* the socket. The socket will not come up easily, and if it does you will have damaged the motherboard.

A point on installing memory in PS/2s and XT/286s: they use a strange "miniboard" of chips called a SIMM (Single Inline Memory Module). They can't be separated: you must replace the whole 9 chip group. Makes diagnosis easier, but fixing more expensive. (Buy a Compaq.)

6. POWER SUPPLIES AND POWER PROTECTION

The PC doesn't come with batteries included. You plug it into the wall socket and it works. The PC itself does not use line current, as this is 120 volt alternating current. Like most digital devices, it needs fairly low power direct current at 5 and 12 volts.

The power supply doesn't *supply* power – it *converts* it from AC to DC. It is a "switching" power supply, which means that it provides more than one voltage to the PC.

6.1. COMPONENTS OF THE POWER SUPPLY

The power supply is the black or silver box in the back of the PC with the large yellow label telling you in five languages not to open the box up, warning you that it is dangerous. Despite the fact that I can only understand a few of the multilingual messages, I'm inclined to take them at their word.

The reason is mainly due to a large capacitor in the power supply. In order to smooth out some power glitches, a capacitor is utilized. The capacitor is a good thing, in general. But it retains power like a battery. Thus, even though the power supply is unplugged, it can still do you some harm. Power supplies cost under \$100. Just replace them if they're faulty. I recommend replacing and not repairing floppies just because it's a pain to repair them: I recommend not repairing power supplies because they can hurt you. (I still recall a shock from a high school power supply project.)

On the side of the power supply is the on/off switch for the computer. Sprouting out of the other side are the power connectors. You will recall P8 and P9 from the discussion of circuit boards. Refer to that discussion for test voltages. The other connectors — two or four of them — are for floppy, tape, or hard drives.

6.2. MAINTENANCE

Good news here: there isn't any maintenance required. The fan and the power switch are the only moving parts. If you suddenly notice that the PC is very quiet, but still operating, then your fan may have died. If that happens, save everything and *shut down immediately!* The computer's own heat can quickly damage or destroy itself if the fan isn't present to dissipate it.

Don't block the vents that the PC uses for cooling. And take compressed air and blow the dust out of the fan now and then. Please remove the power supply first so the dust doesn't immediately settle on the computer.

6.3. UPGRADING THE POWER SUPPLY

IBM PCs and some clones have wimpy 63.5 watt power supplies. This is insufficient for many applications. Memory boards, hard disks, 3278 emulators, and other plug-in enhancements can raise the total power requirements over the 63.5 watt supply's capability. But even if you're only using 50 watts, you should think about upgrading.

A 63.5 watt power supply runs hotter supplying 50 watts than a 135 watt power supply. A bigger power supply means a cooler PC and a longer-lived one. Don't confuse "watts capacity" with "watts used." Sure, a 150 watt bulb uses more power than a 60 watt bulb — here, it's watts used. Both a Volkswagen Beetle and a Porsche 944 can travel at 65 miles per hour, but it seems like real work in the Beetle.

Power supplies are even made which are quieter and/or cooler than the PC. Some example upgrade power supplies are:

- The Silencer 150 and 200: quieter than the XT or AT power supplies (rated at 150 and 200 watts), makes the inside of the PC about 15 degrees cooler.
- Turbo-Cool 150 and 200 are about as noisy as the regular XT and AT supplies, but will cool the inside of the PC by up to 45 degrees.
- The Turbo-Cool 150 is \$159, Turbo-Cool 200 is \$199, the Silencer 150 is \$149 and the Silencer 200 is \$189.

All supplied by: PC Cooling Systems
31510 Mountain Way
Bonsall, CA 92003
(619) 723-9513

Or, you could install an internal "helper" fan to pull more air through the PC, helping to cool it better. Such a device is the ColdBlue from Mandrill Corporation. It cools the inside of the PC up to 25 degrees, which may allow you to run your PC in non-airconditioned environments in the summer. It fits inside the PC case. It costs \$185, and Mandrill is located at:

P.O. Box 33848
San Antonio, TX 78265
(800) 531-5314

Another benefit from an upgraded power supply is additional floppy connectors. The XT and PC power supplies only offer two connections for drives. Many third party supplies have four connections. You could then have two half height drives, a hard disk (half height), and a tape drive, as well as the power to drive them.

The XT and PC use the same case type for the power supply. PC and XT replacement power supplies are about 135 - 150 watts. The AT power supply case is a bit larger, and replacements are in the 190 - 200 watt range. You probably would not replace your AT power supply for reasons of cooling or power except in unusual circumstances. You *may*, however, replace it because the fan is so noisy.

Remember the discussion of the power supply lines in the chapter on circuit boards? You may have noticed the first line on P8, called "Power Good." This is a digital signal enabled by the power supply once it views itself as warmed up and ready. A flaky "Power Good" leads the computer to issue a long beep or short beeps or generally unusual noises. Some inexpensive mail order power supplies cause computers to emit a loud or long beep, then settle down to good service. I have experienced this myself, and can only account for the beeps if the power supply doesn't wait quite long enough to first enable "Power Good." A little initial up and down activity on the line would induce the clock to issue a RESET command to the PC.

6.4. TROUBLESHOOTING THE POWER SUPPLY

You turn the computer on and nothing happens at all. It is plugged in, so it's not that — where next?

6.4.1. The Power Supply Troubleshooting Trail

First, check the outlet. The outlet should be providing between 104 and 130 volts AC current. Just set the VOM to read AC current and put one lead in each hole of the outlet.

Second, check the cables. The cables should be in place on the system board.

Third, is power getting to the power supply? The fan gets it first, so if it isn't turning then the power supply isn't getting power. When the power supply is first turned on, the speaker emits a low click.

Fourth, test the power supply outputs. Acceptable ranges are in the "Circuit Boards" chapter.

6.4.2. Replacing A Power Supply

If you suspect the power supply, replace it. It's simple.

1. On the back of the PC, you will see four screws bolting the power supply to the chassis. Remove these.
2. Disconnect the system board (P8 and P9) and the drives. Draw a picture and make notes of what connects to what. Note wire colors.
3. Slide the power supply forward just a bit — it is hooked to the chassis from below. It will now lift out.
4. Install the new power supply by reversing the procedure.
5. To be extra careful, strip the PC down to the minimum circuit boards. Then power up and do whatever diagnostics you use.

6.5. PROTECTING THE POWER SUPPLY

You can control a lot of things in your environment, but you have little control over one aspect of the PC environment: the power delivered by the electric company. For various reasons, it doesn't come out clean and regular like it's supposed to. This has led to a business catering to those persons requiring laundered power. Devices which perform this function are called surge suppressors, spike isolators, and uninterruptible power supplies (UPSes). They're not all created equally. The rest of this chapter examines these devices and what to look for in one.

6.5.1. Surge Suppressors And Spike Isolators

When some outside force causes your power line to deliver more voltage than it is supposed to, this is called an *overvoltage condition*. Such conditions are, in general, dangerous to the computer.

The physics of it is this: the heart of the computer resides in its chips. The chips are a specially designed crystal. Crystals are highly structured molecules: many of them would be happier in a less structured environment. Applying electronic and heat energy to the crystals allows this breakdown in organization to occur. One spike might not do it, but it leaves damage which is cumulative. Even small spike damage is cumulative.

Damage is proportional to energy. Energy is voltage times current times time.

Brief overvoltages under a millisecond in length are called *spikes*. Longer ones — milliseconds to seconds — are called *surges*.

There are two problems which suppressors face. They must be fast and able to handle large amounts of power. The spike may come and go in 100 microseconds, and the isolator must be able to respond that quickly. (That's why fuses are no good — fuses take seconds to respond.) When it comes, it may be very large. Power line spikes over 15,000 volts have been measured on public utility feeds. This time until redirection starts is called *clamping time*.

The problem is that, in general, fast switching devices are delicate. Devices able to drain more power take longer to act. For example, we have already discussed how hard it is to keep memory chips — a device which switches in under a tenth of a microsecond — from coming to static harm. A very slow switching device like a relay may be powered by very high voltage.

The idea with suppression devices is that once they see a large surge coming they redirect it out to the electric ground — kind of like opening the flood gates. The most common redirection device is called a Metal Oxide Varistor, a MOV. It is an impassable barrier between the supply voltage and protective ground *until* the voltage reaches a certain level. *Gas discharge tubes* and *pellet arrestors* are slower but beefier devices. *Coaxial arrestors* fit somewhere in the middle.

The best suppressors use several lines of defense: MOVs, coax arrestors, and gas discharge tubes, for example.

Of course, an overzealous surge suppressor can redirect *too much power* for too long, and create a worse surge of its own.

Another important question is, *what* voltage level triggers the surge suppressor? They're not waiting for 120.00001 volts to get going: some will pass 1000 volts before calling in the Marines. By then, your PC is toast.

PC magazine did tests of surge suppression devices in its 27 May 1986 issue. They created spikes and measured how much of the spike was allowed through. Some suppressors emitted smoke and flames when subjected to a real surge. Others died quietly, not informing the owner that they no longer protect the PC (they still pass electricity, so there's no way to know.) The most impressive performance was turned in by PTI Industries' Datashield 85, a six outlet suppressor that lets no more than 290 volts through. Some of the more commonly known ones, like Curtis Manufacturing's Diamond six-outlet suppressor passed 1,440 volts. Two good products are:

The Datashield 85 (\$90)
(408) 429-6881

LEA Dynatech Surge Eliminator (\$109)
(213) 944-0916

Summarizing: with surge suppressors/spike isolators, you are interested in:

- clamping time
- maximum voltage
- enough power to withstand an occasional large surge
- multiple outlets
- cost under \$120

6.5.2 Backup Power Supplies

In addition to protection from short power irregularities, you may need backup power. I have lived in a number of places in the Northeast where summer lightning storms will kill the power for just a second — enough to erase your memory and make the digital clocks blink. These are in the range of \$350 to \$1200 and up.

There are two types, *standby power supplies* (SPS) and *uninterruptible power supplies* (UPS). SPSes charge the batteries while watching the current level. If the power drops, they activate themselves and supply power until their batteries run down. A fast power switch must occur here, and it's important to find out what the switching time is. 10 ms or under is fine. 14 ms, in my experience, is not fast enough.

A UPS constantly runs power from the line current to a battery, then from the battery to the PC. This is superior to an SPS because there is no switching time involved. Also, this means that any surges affect the battery charging mechanism, not the computer. A UPS is, then, a surge suppressor also.

A UPS or SPS must convert DC current from a battery to AC for the PC. AC is supposed to look like a sine wave. Cheaper UPS and SPS models produce square waves. Square waves are bad because they include high frequency harmonics which can appear as EMI or RFI to the computer. Also, some peripherals (printers in particular) can't handle square wave AC. So, when examining UPSes, ask whether they use square wave or sine wave. Some produce a pseudo-sine wave. It has the "stairstep" look of a square wave, but not as much harmonic problems.

Ordinarily, the purpose of a UPS is to allow you enough time to save whatever you're doing and shut down gracefully. If you are in an area where the power may disappear for hours, and may do it regularly, then you should look for the ability to attach external batteries to the UPS so that you can run the PC for longer periods. Some allow you to attach inexpensive 12 volt batteries such as those found in cars.

Features with UPSes:

- how much power (watts) is provided?
- for how long?
- how low can line voltage go before the UPS is activated?
- what voltage is required to bring the line current back to the PC?
- sine or square wave?
- can you attach external batteries? Are they regular automobile type?
- two outlets? (one for the computer, one for the monitor)

7. HARD DISK DRIVE MAINTENANCE AND FAILURE RECOVERY

Hard disks have become an absolute necessity. Some programs require them, and most are enhanced by them. But they're a common source of failure, and so require some maintenance attention.

The bad news is, you can't directly fix many disk problems. Even replacing a controller may require reformatting the disk. The good news is, there are many programs to help you monitor your disk's health, tune up its speed, and to recover from disaster.

7.1. DISK TERMINOLOGY

First, some terminology. Disk technology has gone through a lot of changes in just a few decades. There are a lot of options for disks today, and a wide variety in price. What do the parameters mentioned in the ads mean? We'll briefly review them.

7.1.1. Cylinders, Heads, Platters, Tracks And Sectors

A hard disk drive contains rigid metal disks, called *platters*, stacked up inside an air-filtered enclosure. The original XTs had two platters. Other drives may have more. A floppy is like one "platter," but it's not rigid.

Like a floppy disk, the hard disk has a electromagnetic read/write *head* for each side of each platter. Thus, on the original XT drive, there were four heads: 2 heads/platter times 2 platters. As the disk drive is sealed, you cannot align or clean these heads.


Each side of each platter (sometimes called a *surface*) is divided into concentric tracks, like a floppy. A floppy typically has 40 tracks, but hard disks start at 305 tracks and go up from there. The platter is 5-1/4" in diameter like a floppy, so obviously the tracks are squeezed closer in a hard disk than a floppy. Each surface is then figuratively divided up with cuts, like a pie. A floppy, recall, divides each track into 8 or 9 sectors. Hard disks divide each track into 17 sectors.

As with a floppy, each sector contains 512 bytes of data.


Reviewing, on an XT 10 megabyte hard disk,

- 512 bytes goes on each sector,
- there are 17 sectors to a track,
- there are 305 tracks on each surface,
- there are two surfaces (top and bottom) on each platter,
- and there are 2 platters.

$512 \times 17 \times 305 \times 2 \times 2 = 10,618,880$ bytes.




40 TRACK FLOPPY DISK STRUCTURE



40 TRACK FLOPPY DISK STRUCTURE

A TWO PLATTER HARD DISK

A FLOPPY HAS ONE "PLATTER" WITH JUST TWO SURFACES. A HARD DISK MAY HAVE MULTIPLE PLATTERS. BELOW IS PICTURED A TWO-PLATTER HARD DISK.



COLLECTIVELY, ALL TRACK 0'S TOGETHER ARE CALLED CYLINDER 0.

As you examine other disks, you will find that the 512 bytes/sector and 17 sectors/track are fairly constant. The number of tracks and the number of platters vary. Most manufacturers do not report the number of platters, but instead the number of heads, which equals the number of surfaces. The XT has 4 heads.

All 4 heads are attached to a bracket and the same stepper motor. That means that when head 0 (as usual, we count 0 to 3, not 1 to 4) is positioned over track 142 on surface 0, head 3 is also positioned over track 142 on surface 3. Disk heads cannot be independently positioned.

The process of reading a sector involves two steps. First, move the read/write head over the desired track. Then wait for the disk to rotate so that the desired sector is under the head, and read. In general, moving the head takes the most time. This means that we'd most quickly read a file whose sectors were on the same track, and *whose tracks lay above one another*. Thus, if you need track 271, surface 1 for the first 17 sectors of your data on some file, then it would be nice to be able to get track 271, surface 0, 2 and 3 for the rest of your data. Then these data — 512 bytes/sector x 17 sectors/track x 4 tracks = 34,816 bytes — could be read without moving the read/write head. The collection of a given track on all surfaces is called a *cylinder*. Most manufacturers do not report the number of tracks: they report the number of cylinders.

Let's take another example. I am using a Rodime disk drive. It has 639 cylinders and 6 heads. What is its capacity?

We'll assume that, as with 99% of the drives, it has 512 bytes per sector and 17 sectors/track.

512 bytes/sector x 17 sectors/track x 639 tracks/surface x 6 surfaces = 33,371,136 bytes. Recall that a megabyte is not a million bytes, but 1,048,576 bytes, and you see that this is a 31.8 MB drive.

You can find this information on any disk by running the INFO.EXE program on your program disk. A sample run follows. It originally appeared in PC Tech Journal.

Sample Run Of INFO. EXE

*** Information for Disk C: ***

Volume DOS32INT3 created Feb 13, 1987 6:40p
 OEM name : IBM 3.2 Media descriptor (hex); f8
 Volume has 6 Surfaces, 639 Tracks with 17 Sectors/Track
 Sector size is 512 bytes. FAT entries are 16 bits
 Cluster size is 2048 bytes (4 sectors)

Usage:	Sectors	Bytes	Clusters
Sectors Not Available	0	0	0
DOS Boot Area	1	512	
File Allocation Table	128	65536	
Root Directory	32	16384	
Files & Subdirectories	20892	10696704	5223
Locked Out	0	0	0
Available	44108	22583296	11027
TOTAL	65161	33362432	16250

The disk is 32% full

7.1.2. Disk Performance I: Seek Times, And Latency Periods

The hard disk drive is somewhat like a phonograph. The platter itself is like the record, and the read/write head is like the phonograph stylus. To play a selection, you do two things:

1. Position the stylus arm to the beginning of the song
2. Wait for the music on the track to come around

Obviously, 1 takes a lot longer than 2. Almost instantaneously after positioning the stylus, the music is found.

Disk reads and writes work the same way. The head must first be positioned over the track, then it waits for the desired sector to come up.

The two components are called *seek time* and *latency period*. Seek time is the time required for the head to position over a track. Latency period is how long it takes for the desired sector to move under the head.

The seek time is usually the larger wait. It varies according to how many tracks must be traversed. A seek from one track to the next track is usually quick — 5 to 16 ms, but most seeks aren't so convenient. A common measure of an "average" seek is the time required to travel one third of the way across the disk. This is the one used in the benchmark program most commonly used in the PC world, the CORETEST program (more on this later). You might wonder, "why not half-way across the disk, rather than one-third?" The reason is that most accesses are short seeks — just a few tracks.

Following are seek times for a range of computer/disk combinations:

Toshiba 3100, 10MB	175 ms
XT original 10MB	100 ms
Rodime 33 meg in XT	57 ms
AT original 20 MB	30 ms
Priam 60 MB	22 ms
Maxtor 140 MB in AT	19 ms

Once a head is positioned over a track, it still takes time (unless you're lucky) for the sector desired to rotate under the head. If you're lucky, it's already there. If you're really unlucky, you just missed it and will have to wait an entire revolution. This time is the *latency period*. A common number cited is average latency period. This makes the simple assumption that, on average, the disk must make a half-revolution to get to your sector. Recall that the disk rotates at 3600 rpm. One half revolution then takes $1/7200$ of a minute = $60/7200$ second = 8.33 milliseconds (ms). This contributes to the amount of time that the system must wait for service.

The sum of the average seek time and the latency period is called the *access time*, and is oft-quoted in product advertisements.

7.1.3. Disk Performance II: Data Transfer Rates and Interleave Factors

Once a disk has found the desired data, how fast can it transfer it to the PC? This is called *disk transfer rate*.

The hard disk rotates constantly at 3600 rpm (compare to 300 rpm for a floppy). Even when you are not accessing data, the drive spins. Floppies, of course, only run the motor when data is required. Hard drives must run continuously because the head does not actually touch the platter: it flies above it on a cushion of air. When the disk is spinning, the cushion of air is present. Another reason for the drive to run constantly is to speed up access. Floppies wait a half second for the motor to get up to speed each time you read or write data.

There's a downside to the fast rotation. Say DOS is going to read all 17 sectors on a given track. We move the head and wait for sector 1 to come under the head. We read it and pass it to DOS, which processes it and then asks for the second sector. But while DOS was processing, the drive kept spinning. We might be over the third sector *after* the one we just read. (3600 rpm is fast!) This leads to the notion of *interleave factor*

If the sectors are arranged in order physically around the disk, like so:

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 (wraps around to 1...)

then the disk is inefficiently accessed. We read sector 1, and perhaps the head is over sector 6 by the time DOS wants sector 2, so we've got to wait for another go-round. By the time we want sector 3, the head is over sector 7 and so on. How can this be fixed? Assume that we will end up skipping three sectors between one read and the next. We'd like the disk to look like, then

1 (skip 3) 2 (skip 3) 3 (skip 3)

This is easily done. The sectors are not laid out in consecutive order. When the disk is physically formatted (see later section on this), the first sector is laid down, then the third sector after sector 1 is designated sector 2. The third sector after sector 2 is designated sector 3, and so on. This is called *interleaving*, and in this case would be called a 1:3 interleave factor. The resulting disk would have sectors in the following order:

1 7 13 2 8 14 3 9 15 4 10 16 5 11 17 6 12 (wraps around to 1...)

Notice here the benefit of a 17 sector disk: just about any interleave value works fine. (Try doing a 1:4 interleave on a 16 sector disk).

The original XT required a 1:5 or 1:6 interleave. A really fast computer could use 1:1 (noninterleaved) disk. The computer that I am working on (a 10 Mhz XT clone) has optimal interleave of 1:3. An interleave which is too small or too large slows down your disk unnecessarily. When you do a physical format, discussed later in this chapter, on your disks, spend some time and locate the optimal interleave factor. The defaults are often not the best value. *On an IBM XT, re-formatting to a 1:5 interleave will improve disk throughput by 20%.*

See pages 72 and 73.

A program, HOPTIMUM, packaged in the Kolod Hard Disk Utilities, will measure the optimum interleave factor for your disk. You can get these utilities (we will discuss them more later) from Paul Mace Software:

Paul Mace Software
123 N. First Street
Ashland, OR 97520
(503) 488-0224

Typical HOPTIMUM output is on page 71.

SAMPLE HOPTIMUM OUTPUT

C:\MACE>hoptimum

HOPTIMUM - Hard disk interleave OPTIMIZER Program - Version 1.0.3

Copyright (c) 1986 Kolod Research, Inc.

Box 68 * Glenview, Il. 60025 * 312-291-1586

Licensed Program Materials

Drive - C - will be formatted and tested as:

read/write heads: 6
cylinders: 639
sectors: 17

Formatting parameters will be:


beginning surface: 0 (surfaces start internally at 0)
ending surface: 5
beginning cylinder: 628 (cylinders start internally at 0)
ending cylinder: 638
sector interleaves: 1 through 7
machine type: PC/XT (or compatible)

Do you wish to change any of the above
formatting parameters? (Y/N): n

...more output...
Summary is as follows...


Interleave	Total access time (rounded to nearest sec.)	Total clock ticks (normally 18.2 / sec)
1	20	364
2	20	365
3	4	65
4	5	86
5	6	106
6	7	126
7	8	146

Program has terminated normally



17 Sector Disk

1 : 6 Interleave



17 Sector Disk

1 : 3 Interleave

7.1.4. Testing Disk Performance: CORETEST

CORE International sells some of the fastest disks around. It's no good being number one (or close to it) if no one knows it. So they wrote a benchmark program and put in the public domain. It's called (appropriately) CORETEST.EXE, and is on your program disk. To run it, just type **CORETEST**. Below is a sample run. Here's how to interpret it.

The two interesting numbers are "Data Transfer Rate" and "Average Seek Time." Average Seek Time can be added to 8.3 ms to yield average access time. In the case of the sample run, $57.6 + 8.3 = 65.9$ ms access time. *Seek times are unaffected by interleave factor!* To tune interleave factor, examine Data Transfer Rate. Do a physical format for a given interleave, then create the partition with FDISK, then do a DOS format, then run CORETEST. Note the rate. Then do a physical format with another interleave factor and repeat the rest of the process, noting the Data Transfer Rate. Try this for interleave factors 1 to 7. Choose the interleave which lead to the greatest Data Transfer Rate. Or buy the Mace HTFORMAT program — it optimizes the interleave factor for you.

Sample CORETEST Screen

CORE Disk Performance Test Program Version 2.6
Copyright © CORE International, Inc. 1986

Seek Times		Hard Disk 0		KBytes Read
80+	Size: 33.4 MB		Heads: 6	2048 +
75	Cyls: 639		Sects: 17	1920
70	Data: 663KB	HD0	Time: 4.0 secs	1792
65		Data Transfer Rate: 164.3 KB/sec		1664
60		Average Seek Time: 57.6 ms (639 cyls)		1536
55		Track-Track Seek: 8.6 ms		1408
50		Performance Index: 1.933		1280
45				1152
40				1024
35				896
30				768
25				640
20				512
15				384
10				256
5				128
0	HD0		HD0	0

Transfer Block Size: 51KB

7.1.5. How Much Faster Is A Hard Disk Than A Floppy Disk?

What am I buying besides greater storage size? Several excellent articles on hard disk performance appeared in a special issue of *Byte* devoted to the IBM PC in the fall of 1985. The following numbers come from "Fixed Disks and IBM PC AT."

Drive Type	Effective Data Rate (bits/second)
360K Floppy, XT	250,000
360K Floppy, AT	300,000
1.2M Floppy, AT	500,000
Hard Disk 1:6 (XT)	833,336
Hard Disk 1:3 (AT)	1,666,664
Hard Disk 1:2 (AT)	2,500,000

In reality, a hard disk will be 5 to 10 times faster than a 5-1/4" floppy.

7.2. COMPONENTS: THE CONTROLLER AND THE DRIVE

The hard disk subsystem consists of the sealed drive itself, and a controller board which plugs into a slot in the PC expansion bus. Controllers are available which act as both floppy and hard disk controller, saving you a slot. Some firms have developed "hard cards" which mount a narrow drive on a short controller board to yield a plug-in hard disk and controller all in one. We will consider the two components separately.

7.3. THE CONTROLLER

The controller is a printed circuit board that generally contains some VLSI chips on it, but not a microprocessor.

7.3.1. The XT Controller Vs The AT Controller

Some controllers talk with the computer via DMA (XT types), and others use an odd IRQ approach.

In the IRQ approach (used in the AT), the controller first fills a 512 byte buffer on the controller board itself. Then a hardware interrupt type 14 is issued, and the 80286 CPU reads the data from the buffer and stores it in memory. This sounds counter-intuitive, as the whole idea of DMA is to be able to transfer data into memory without the CPU slowing you down. The reason that it works is due to the design of the interface between the controller's buffer and the 80286. Most controller designs spoon out data in 8 bit bytes. The AT controller hands data to the CPU in 16 bit chunks. A second reason for the success is two commands that the 80286 has that the earlier chips don't have, which allow fast block transfers from an I/O port. The raw transfer rate for these transfers can be as fast as *16 million bits/second*.

7.3.2. MFM And RLL Controllers

Another way to get better speed, *and more storage*, is to use a kind of hard disk controller called an RLL (Run Length Limited) controller.

For years, computer manufacturers have tried to squeeze as much data on a given disk drive as was possible. Consider, for instance, that the original 5-1/4 inch floppies stored about 90K of data: now some store 1.2 or 3.3 MB. How is this done?

Basically, data is stored on a magnetic medium by encoding what are called *flux reversals* onto the magnetic medium. A reversal means either "negative to positive" or "positive to negative." The easiest approach, it seems, would be to encode something like

0 = no flux reversal
1 = flux reversal

But the problem would be that if we got a long string of zeroes, then there would be a long period of no reversals. We *could* time how long it is between flux reversals and deduce the number of zeroes, but that makes some heroic assumptions about available timing devices. Unless your timing was very accurate, you'd probably drop a zero or two.

The more "fault tolerant" approach is to include clocking bits with the data. One simple approach is to code a 1 as two reversals, then code a 0 as a reversal followed by no reversal:

1 0 0 0 1 1 becomes
RRRNRNRNR (R=reversal, N=no reversal)

This is, as you can see, effective but wasteful. To see this, you need one more piece of the puzzle: *you can fit more "reversals" and "no reversals" on a disk if the number of "reversals" is relatively small.* The scheme pictured above, as you can see, includes a lot of reversals. It does eliminate the "long run of zeroes" problem, however. The minimum length run of zeroes is 0 — that is, sometimes two reversals occur next to each other. The maximum length is 1 — you never have more than 1 zero in a run. We can say that this is a (0,1) run length limited scheme. This is an old encoding scheme, called "Frequency Modulation," or "FM." It's not used much anymore.

FM lead to MFM — Modified FM. MFM was a bit more daring, creating minimum zero runs of 1 zero and maximum runs of 3 — a (1,3) run length limited scheme. IBM uses this for floppy and most hard disk encoding.

Other approaches exist. Apple uses something called Group Coded Recording — GCR. It stores data on the disk in groups of four bits. There are 16 four bit combinations. Some — like 0000 — have no ones in them. With GCR, you take all *five* bit combinations (there are 32) and remove the ones with more than two zeroes in a row. There are 17 of them. Create a table corresponding the 16 four bit combinations to the 17 legal five bit combinations. The 17th combination is not used for data. The original four bit combinations are then represented on the disk as five bit groups.

Recently, a new kind of controller has arrived - 2,7 RLL. It pushes the density of data on a hard disk a little further — 50% further. Common RLL configurations combine the Seagate ST4096 (a 21 MB drive) and an RLL controller to yield a 30 MB system. (Seagate also makes real 30 MB disks, like the ST238 — 31.5 MB). *Be careful with these.* Experts say they're a proven technology, but I've seen a higher proportion of disk problems with RLL controllers. As the RLL technology pushes the disk a bit closer to its limits, many disks will work fine on MFM but not RLL. The Seagate ST4096 is fine, but don't try just any disk (my Rodime, for example, can't use RLL).

A final benefit from RLL — since the data is stored in a denser fashion, it is transferred off the disk faster. When RLL works, it makes your disk *faster* and *bigger*.

7.4. INTERFACES

How does the controller talk to the drive? Obviously, an interface like RS232C isn't fast enough. A common interface is called ST506. An ST506 interface is 5 million bits per second, a parallel interface.

Another interface not used in the PC world much now, but on the horizon is SCSI (pronounced — I kid you not — "scuzzy"), Small Computer Systems Interface. Probably the best selling SCSI computer is the Macintosh Plus, and no doubt soon the Mac II. It transfers data at 10 Mbps. The Bernoulli Box uses SCSI also, on the PC.

A controller board is commonly connected to a drive with two ribbon cables. The larger has 34 lines, the smaller 20. Many controllers have one 34 pin connector and two 20s. The second 20 is for connecting a second drive.

When connecting the ribbon cables, *remember* that the cable will be gray or white with one edge red. The red wire goes to pin 1. As I said before, I made the mistake of reverse connecting a disk once (*once*). While waiting for the PC to boot up so that I could test the disk, I left the room for a second. Returning, I was greeted with "making toast, Mark?" (Co-workers can be so amusing at times like this.)

7.4.1. Informing The Controller Of Drive Type

Controllers and disks aren't manufactured together. Controller manufacturers, like Xebec, Data Technology Corporation (DTC) and Western Digital build controllers to be generic. For example, when I first received my 33 MB Rodime drive, it was supplied with a Xebec controller. I formatted the drive and used it for a while, finding it a bit slow. A CORETEST showed an access time of 639 ms — substandard or, as the tech support person said, "*sub* - substandard." I got a DTC 5150X controller and connected it — again, standard cables and connectors. It was built to handle a variety of disks. Disk type 4 was 639 cylinders, 6 heads — the Rodime. I selected disk type 4 with a set of jumpers. (On an AT controller, recall, you do it with the SETUP program.) The drive formatted without trouble and yielded a seek time of 57.5 ms.

7.5. THE DRIVE: THE "WINCHESTER" SEALED DRIVE

On some mainframe disk drives, the disk is in a separate "disk pack" which mounts atop the drive motor and control circuitry. IBM changed this years ago with the "fixed" drive — a drive sealed in an air filtered compartment. It could not be mounted or dismounted. (That's why IBM calls the hard disk the "fixed disk drive." I know of no IBM intentions to offer a mountable drive for the PC, but you never know.) The model number was the 3030 drive. The model number reminded technicians of the Winchester .3030 rifle, and the drive became known as a "Winchester" drive. The name has stuck through model numbers.

Winchesters for microcomputers have until recently been based on 5-1/4" platters. It is interesting that whether the drive is 5 MB or 180 MB it is contained in the same size housing. You can't look at a drive and detect its capacity by its physical size.

New 3-1/2" platter technology is now starting to appear, mainly in the form of so-called "hard cards" and in portable PCs like the Toshiba 3100.

7.5.1. Hard Cards

A "hard card" is just a printed circuit board with the controller and drive mounted on the same board. The drive must obviously be fairly small. This has led to some problems with these devices. I must frankly say that I'm at a loss to understand their popularity. They are less reliable, slower, and more expensive than their discrete-component counterparts. One argument leveled in their favor is the lower power drain. As far as I'm concerned, any significant expansion in the PC requires an inexpensive upgrade power supply, so get one anyway. It's still cheaper.

7.5.2. You Can't Fix It, So You Must Protect It

There's not much to be done in the way of regular maintenance on the sealed drive: you can't clean the heads, you can't align them. If the drive itself is fried, then you can cut it open and leave it on your bookshelf as a conversation piece.

What kills hard disks?

- smoke (the filters aren't perfect)
- vibration

Smoke is straightforward, and has been discussed before. When a smoke particle sits on the disk, the head whips around and crashes right into it. On a 5-1/4" platter rotating at 3600 rpm, a particle on the outermost track would travel at 112 miles per hour. The outermost track is, by the way, the most important track, as it contains the directory information for the rest of the disk.

Vibration is a problem because of the way in which the head travels above the disk platter. Recall that the head travels above the disk on a cushion of air. The distance between the platter and the head (called the *head gap*) is fairly small. Any bouncing around, then, can cause the head to literally crash into the platter below. This plows a small furrow in the platter, and grinds down the head. Do this enough, and you'll lose the data in the furrow and possibly the head.

A similar thing happens whenever you turn off power from the hard disk. Deprived of its cushion of air, it crash-lands onto the platter below. If you always have the head positioned over the same spot when you power down, you'll eventually create a new bad sector there.

A way to reduce the damage from this is to position the head to a location outside of the data area, *then* power down. This can be done because the actual number of tracks that you have is greater than the number you use by a few. For example, on the 305 cylinder XT, you can position the head over cylinder 306, then power down. This is called *parking the heads*.

Another way to protect the data is seen in the way that magnetic material is put on the platters. The cheaper and less effective approach is to simply *coat* the platters — in effect, to glue rust onto the platters. A better and more expensive approach is to *plate* the magnetic oxide onto the platters. Plated media is superior, and more expensive.

7.5.3. Band Steppers And Voice Coils

Less expensive, slower drives use a technology called *band stepper* to position the head. XT disks are usually band steppers. Faster, more expensive drives use *voice coil* technology: it's faster and more precise. A side benefit of voice coil is that the heads are held against the innermost, unused tracks when power is off. You must apply power in order to move the heads out over the data. This means that voice coil drives are self parking.

Band steppers, on the other hand, are not self parking. IBM included a program called SHIPDISK on the XT diagnostics. The notion is that when you're going to subject the disk to some jarring, like when you're moving the XT, you should first boot the Diagnostics disk and use SHIPDISK to position the head, then turn the system off.

SHIPDISK works fine *when you use it this way*: boot off the Diagnostics disk and then run the program. But a lot of people decided to just copy SHIPDISK off the diagnostics disk onto their hard disk, and run SHIPDISK before powering down for the night. This is a bad move, as SHIPDISK presumes the presence of some software routines which only exist on the special version of DOS included on the Diagnostics disk. Run under regular DOS, SHIPDISK can destroy data.

What, then, to do? Two possibilities. First, you could leave your PC on all of the time, as has been discussed earlier. Second, use one of the park programs on your program disk: DPARK-PC.COM for the PC, DPARK-XT.COM for the XT, and DPARK-AT.COM for the AT.

If you work in an environment with a lot of vibration, don't hesitate to park the heads whenever you're not using the PC, even if it is still turned on. See the program disk for AT disk parking programs for use in this unusual case.

7.5.4. Squeaky Drives

Many hard disks make high pitched squeals and squeaks. How to handle this? First, ensure that the disk is fully secured, and either lying flat or on its side. A slightly askew disk will whine. *By the way* — it's perfectly acceptable to store a hard disk on its side, but you'll get better results if you also *format* the drive in this position.

The second approach involves lubrication. Remove the drive. On its bottom there is generally a circuit board. The board may have a cutout in the middle: if not, remove the board. You'll see a small bearing. Using WD40, put *just a drop* of lubricant on the bearing.

7.6. THE BEST INSURANCE: GOOD BACKUPS

There are two kinds of hard disk users: the ones who have had a disk failure and lost data that wasn't backed up, and the one's who are going to⁷. You seek to stay in the first category as long as possible.

I can't stress this strongly enough: **at the first sign of unusual behavior, back up the entire disk.**

Backup approaches fall into three categories:

- software-only
- special hardware
- PBB (Prayer-Based Backup: hope it'll never happen)

Informal surveys indicate that about 20% of the corporate users employ special hardware and the other 80% claim to use software-based schemes. Many of the 80% use a hybrid software/PBB system.

⁷Actually, I first heard this as "there are two kinds of pilots: the ones who have landed and forgot to lower their landing gear, and..."

7.6.1. Software Based Backup

These programs assist the user in transferring data onto floppies. Their main features are:

- speed
- maintain indices of backed up data
- use of the DOS archive directory bit
- splitting large files over several floppies
- file compression
- error correcting data formats

Despite its difficulty and bugs, the DOS BACKUP program is probably the most used software alternative. It is not particularly fast, and it stores data in a nonstandard format, but it comes free with DOS. It will break up files which are too large to fit on a floppy to several floppies, and manipulates the DOS archive bit. No compression or error-correcting is used.

Fifth Generation Systems' Fastback product is well-known and widely used. It is much faster than BACKUP, and formats disks on the fly. It uses error correction but no compression, and stores data in a nonstandard fashion.

Similar programs are offered by other vendors. But here's another approach. Under DOS 3.2, a new command — XCOPY — appeared. XCOPY has several features which recommend it for disk backup:

- it is faster than COPY
- you can tell it to prompt you for filenames
- it operates on the archive bit
- the resulting files are in DOS format

The main drawback is that XCOPY is not smart enough to break up large files over several disks. XCOPY is smart enough to only copy the files which have changed since the last backup, and to search subdirectories (with COPY, you must specify each subdirectory). It will also create subdirectories on the target disk. From the root directory, just issue the command

`XCOPY C: A: /V/M/E/S`

XCOPY will then fill up a floppy (it must be formatted first). Insert a new floppy and issue the command again. Continue until there's nothing left to copy. Then use BACKUP for the files too large to fit on one floppy.

No matter what system you use, inserting and removing floppies is a large pain. If backups are difficult, they won't get done. That's why hardware backup solutions are superior, and worth the money.

7.6.2. Backup Hardware

There are several backup hardware options:

- tape drives with special controllers
- tapes based on the floppy controller
- removable hard disks, like the Bernoulli Box
- LANs
- A second Winchester

Tapes cost from \$450 up. For an office of XT's, a good solution would be to get a portable tape drive which can plug into the back of the floppy disk controller (you wondered what that 36 pin port was good for, eh?). Then one \$700 tape drive can be shared among many computers. Irwin Magnetics makes a 20 MB drive that plugs into the floppy controller. A mail order company, Alphanumeric International (213) 921-8689 offers it for \$540.

Bernoullis are reliable, fast devices which double as an extra hard drive (they are just as fast) and backup device. Because there are two drives, backing up a Bernoulli cartridge is simple. To use a Bernoulli to back up a hard disk, Iomega includes backup software, or you can use the XCOPY approach outlined above.

One writer has suggested that, as hard disk drives have become so inexpensive, it could be cheaper to just install a second hard disk and use it to back up the first. (No, the writer doesn't work for Seagate.)

7.7. DRIVE PREPARATION

When a new hard disk is being installed in a system, three steps are done:

- Physical Formatting
- Partition Creation
- Directory Formatting

7.7.1. Physical Formatting

Physical, or "low level" formatting, is a process whereby 17 sectors are "drawn" on each track, using magnetism as the "ink." Each sector is then filled with ASCII 229 (the o, "sigma" character). It is at this level that the interleave factor is set. This formatting is generally done by the manufacturer.

You will do a physical format for one of three reasons:

1. You experience a large number of disk problems. Try cleaning off the hard disk, do a new physical format, create a partition, and directory format, then reload your files. Do not re-format a drive just to get rid of bad sectors. They'll come back.
2. You wish to experiment with the interleave factor to increase your data transfer rate from your hard disk. This should be done when the disk is first installed, and again whenever your system is speeded up, as with an accelerator board.

3. You suspect a faulty format, as in a case wherein a disk is accidentally formatted as a 20 MB drive when it is actually a 10 MB drive.

Once I tried to use a rental XT under DOS 3.1. The XT was delivered to me with the disk formatted under 2.1, so I tried to re-format it. The drive ground a lot, and finally told me that it had formatted a 20 MB disk and had found 10 MB of bad sectors. It was, of course, a 10 MB disk in reality which had gotten an odd formatting. I ran a physical format, then created the partition and did the DOS format, and there was no further trouble.

How you physically format is a controller-specific question. Some controllers (some of the XT variety) have the format program right on the controller ROM. Others have a separate program which must be loaded off disk. The AT controllers require the AT Advanced Diagnostics to re-physically format, or buy one of several third-party format programs.

Many controllers allow you to re-format with a built in program found at address C800:5 or C800:6.

Warning: the procedure which I am about to describe will irrevocably destroy any data on your hard disk drive. Do not do it unless you have backed up the drive in entirety!

Load DEBUG:

```
C>debug
-G =C800:5
```

You should get some kind of prompt, like "What interleave factor would you like to use?" The prompts will depend on the manufacturer: see the documentation that came with the controller. If you don't have documentation, call the manufacturer's tech support line. My DTC controller just asks what interleave factor to use, then goes to it. Western Digital controllers want to know a lot of information about cylinders and such.

Make sure that you know how to do physical formats. If you are responsible for several or many machines, get the instructions and try it out once. The reason for this: you want to see all diagnostic procedures on a properly functioning system. Then you know how to spot troubles.

7.7.2. Partitioning

In order to accommodate multiple operating systems, DOS allows you to create multiple *partitions* with the program FDISK. Suppose you have a 30 MB drive and want to run both DOS and XENIX. Using FDISK, you can create a (say) 20 MB partition for DOS, and a 10 MB partition for XENIX. FDISK is menu-driven, and I'm not going to spend time explaining how to use it here.

What you are concerned about is, *what happens when the partition information is destroyed?* This is a real problem: DOS will refuse to even recognize that a disk exists. A software company, Paul Mace Software, offers some utilities to help:

Mace Utilities and HTFORMAT/HTEST (\$99 a piece)
Paul Mace Software
123 N. First Street
Ashland, OR 97520
(503) 488-0224

We will talk more about the Mace Utilities in the next sections.

As of DOS 3.3, FDISK has another value. DOS (and, for that matter, OS/2) can't handle a disk larger than 32 MB. Previously, this meant that a 40 MB drive wasted 8 MB, or required a device driver to access the latter 8 MB.

Under FDISK 3.3, you can divide the disk into a drive C: with 32 MB and a drive D: with 8 MB. You can create as many volumes as you like, all smaller than 32 MB.

7.7.3. DOS Formatting

Finally, you run the DOS FORMAT program. It does not actually overwrite sectors and physically format hard disks. It just sets up the four areas of a disk:

- The boot record / disk ID area
- The File Allocation Table (FAT)
- The Root directory
- The Data Area

Note that the FORMAT program does not destroy or overwrite data in the data area: it just sets up the boot, FAT, and directory. This implies that *formatted hard disks can be recovered!* See the later section on Failure Recovery for more information.

There are, incidentally, a few exceptions to this rule. AT&T and Compaq DOS 2.11 actually torpedo the whole disk. There is no recovery from this.

7.8. FAILURE RECOVERY

The point of this discussion is to understand how to recover from disk failures. There are several kinds of failures:

- damaged partition
- damaged boot record
- damaged FAT
- loss of data in data area
- accidental erasure of files
- accidental re-formatting of disk

We will provide some background, diagnostic tools, and recovery procedures for these problems.

7.8.1. Damaged Boot Record

To "boot" off a floppy disk, early versions of DOS put some computer code onto the first sector of the disk. The ROM in the PC knew just enough on powerup to load that sector and execute the code found within. This sector was called the "boot record." This sector is, recall, 512 bytes long, like all sectors.

IBM uses the boot record — previously only for bootstrap code — to contain DOS version and disk type info. The 21st byte is important, as it is the "Media Descriptor." The table below lists these descriptors. Another boot record item is "number of sectors on disk." Two bytes are allocated for this. The maximum two byte number is 65,535. As sector sizes are 512 bytes, this leads to the current maximum volume size of $65,535 \times 512 =$ just short of 32 MB.

Media Descriptor Byte Values

Code	Description
FF	8 sector DS floppy (320K)
FE	8 sector SS floppy (160K)
FD	9 sector DS floppy (360K)
FC	9 sector SS floppy (180K)
	(also, strangely, Bernoulli Boxes)
F9	1.2 MB floppy and 3-1/2" floppies
F8	Hard Disk

If your boot record is damaged, you will be completely unable to read the disk, even if you don't need to boot from it. You can reconstruct and repair a boot sector with the help of a special utility like Mace. "Restore Boot Sector" in Mace will either use a backup boot sector (which you must allow Mace to create before the problem exists), or copy a boot sector from an identical disk.

The problem with the boot record is that it must be in a specific physical place on the disk, the first sector. This is the first sector of the first track of the first surface. Thus, if the first sector of a hard disk is bad, the entire disk is unusable.

The boot record contains a few basic pieces of information, so it's a shame that no one makes a simple utility to allow you to build a boot record from scratch. You *can*, however, construct a boot record in DEBUG and write to sector 1. Here's a good way.

Fixing a Damaged Floppy Boot Record

1. Insert a good disk of the same size, density, etc., in A:.
2. Load debug. It's on the DOS supplemental programs diskette, and should be in your DOS subdirectory.
3. Instruct DEBUG to load the boot record off the good disk: type `L 0 0 0 1`
4. Remove the good disk from A: and insert the disk with the bad boot record.
5. Tell DEBUG to overwrite the bad boot record with the good one: type `W 0 0 0 1`
6. Quit out of DEBUG.

Warning: Do not mix boot sectors! Don't try to patch a boot sector on a hard disk with one from a floppy. It won't work. You *could*, if you like, use the boot sector from one hard disk to save another. It's a similar procedure to the above. The idea is to read the boot sector off the good hard disk, write it to a floppy, transport the floppy to the computer with the bad hard disk, read the boot sector off the floppy, and write it back out to the hard disk.

Fixing A Damaged Hard Disk Boot Record

1. On the computer with the good hard disk, insert a blank formatted floppy in drive A:.
2. Load DEBUG.
3. Tell DEBUG to read the hard disk's boot sector and store it on the floppy:
type L 0 2 0 1
type W 0 0 0 1
4. Take the floppy to the computer with a bad boot record on the hard disk. Load DEBUG off a floppy. Insert the floppy with the boot record from the good computer.
5. Instruct DEBUG to read the floppy data and put it on the hard disk:
type L 0 0 0 1
type W 0 2 0 1

This will sometimes work. Sometimes it will not, if other data (like the partition record on the hard disk) is damaged.

PC DOS 3.2 Tip:

Final point on the boot record: the fourth byte through the eleventh byte contains a string with the manufacturer ID (called the "OEM name and version"), like "IBM 3.2" on a PC-DOS 3.2 formatted disk. This was fairly irrelevant until PC-DOS 3.2. As of 3.2, however, the boot record information is absolutely essential. One of the things in the boot record is the location of the FAT. If you have a disk written under a version of DOS that doesn't bother with the location of the FAT, DOS 3.2 could become confused and try to read a FAT where one does not exist. You will get a "sector not found error" type message. This only occurs under PC DOS 3.2 as far as I know. This can be remedied fairly easily. Just format a blank floppy under DOS 3.2, then do the above procedure. Use DEBUG to copy the 3.2 boot record off the blank disk to the other disk.

This problem doesn't always occur, but if it does, try this. Better yet, try it on a *copy* of the disk.

7.8.2. Recovering From A Damaged FAT

The next area is the FAT, the File Allocation Table. The FAT keeps track of which *cluster* is associated with which file. But first, what's a cluster?

A cluster is a group of sectors. It is the minimum space allocated by DOS when DOS gives space to a file. For example, if you create a file which is 1 byte long, you don't take up just a byte on the disk, but instead the minimum allocation — a cluster. Cluster size varies with the disk type. A 1 byte file on a single sided floppy would take up 512 bytes on the disk, as would a 500 byte file. A 1 byte file would take up 4096 bytes on a 10 MB disk drive.

Think of the following analogy. Say you're the guy running the Eastern Shuttle. The Eastern Shuttle promises a seat on a jet going from Washington, DC to New York every hour, without a reservation. Say the DC-9 jets that they use can accommodate 100 people. How many jets must you run this hour? If you've got 1-100 people, run 1 jet. Just one person — the 101st person — forces you to go to a second jet. There are no "half" jets.

Cluster Sizes

Disk Type	Cluster Size (bytes)
Single Sided Floppy	512
Double Sided Floppy	1024
3-1/2" Floppy	1024
1.2 MB Floppy	512
10 MB Disk Drive	4096
16+ MB Disk Drive	2048
Bernoulli Cartridge	4096

Depending on what kind of medium DOS uses, different size "jets" are used. First the floppies appeared, then the 10 MB disk drive, then the 1.2 MB floppies and the 16+ MB drives. Recall that a limitation in the current boot record format limits us to 32 MB drives. The FAT does not limit us, yet: we could, within the current structure of DOS, create a disk with a 16 bit FAT (65,536 entries) and clusters of size 4096, and support a 268 MB drive — *if* the boot record constraint didn't exist.

Anyway, back to the FAT. It is a linked list of clusters, one FAT entry for each cluster. There are two kinds of FATs: 12 bit and 16 bit. A 12 bit FAT can keep track of $2^{12}-1$ or 4095 clusters. It is used for floppy disks and 10 MB hard disks. A 16 bit FAT can keep track of 2^{16} or 65,535 clusters.

Most disk formats keep a second copy of the FAT. I guess it originally was included for fault tolerance or redundancy, but nothing uses it. If your FAT is damaged, you can copy the secondary FAT back to the primary position with DEBUG, or (more simply), Mace will do it. Unerasing programs manipulate the FAT.

7.8.3. Recovering Lost Data Due To Disk Media Failure

As you know, if you have ever formatted a floppy disk, not all sectors are created equal. Sectors that started out reliable can suddenly go south for the winter, taking important data with them. There are some things that can be done in this case, however: you're not helpless.

The initial formatting process tests for bad sectors and ensures that DOS doesn't try to use them. The FAT is where DOS keeps track of bad sectors. If the initial FORMAT locates bad sectors, then it will inform the FAT of them. The FAT entry for the cluster that that sector resides would then indicate that this is a "don't use" cluster.

What about newly-created bad sectors, like when the disk "crash lands" on your surface too many times? These "spontaneous" bad sectors must be reported to the FAT.

There are two parts to this problem. The first one is that the sector probably contained data that you need: you'd like to recover this data. The second is that the sector must be reported to the FAT.

First, recovering the file. You can use a utility to do this, but I prefer to do it by hand, as some of it is just luck. First, COPY the file. In the process, you will probably get a message like one of the following:

Sector Not Found Reading Drive X:
Data Error Reading Drive X:
General Failure Reading Drive X:

Followed by the question: "Abort, Retry, Ignore?"

The data is probably not *dead*, but rather the magnetic field is weakened. (As Monty Python would say, "it's not dead — it's just *resting*.) Retry it to see if the problem can be resolved. Retry several times. When it is obvious that you're not getting anywhere, use Ignore. Here's what Ignore does: the sector is probably not completely garbaged. Only one byte out of the 512 may be damaged. Ignore tells it to copy it, warts and all. In some cases, this may mean an innocuous error. In other cases, it may mean that the data is irretrievably lost.

Now you have recovered as much of the data as is possible. If it has errors, and you want to sew the data back together, *and are stout of heart and long of patience*, you can load up DEBUG and try patching the data. For example, one bad byte in the wrong place can render a dBase or Lotus file useless. You can get a book which describes file formats for the big programs, and use it to try to repair data. The name of the book is File Formats for Popular PC Software, by Jeff Walden, published by Wiley. I stress that this book doesn't tell you how to fix files — just what undamaged ones look like. It is no simple task to rebuild files, but it can be done sometimes with patience and time.

Now update the FAT. There are several programs which do this. One is the DOS program RECOVER. RECOVER does not do the above-mentioned retries and ignores — it just assembles whatever doesn't give it trouble, collects it into a file with a name like FILE0001.REC, and marks the bad sectors on the FAT. The ".REC" file is of no use to you — just erase it (it's in the root directory, by the way).

Another approach is to use the Norton Utilities' DISKTEST program, or the Mace Utilities REMEDY program. These are better programs, as they will scan an entire disk for you, but they don't do the retry/ignore work — you've still got to do it yourself. Mace has a command — DIAGNOSE — which will alert you to bad sectors without trying to update the FAT and kill the sectors. Norton will do the same.

It bears repeating: as soon as your hard disk starts showing spontaneous new bad sectors, back up the whole disk *now*.

7.8.4. Recovering Accidentally Erased Files

Sometimes, carelessness and haste bring a painful cost: accidental erasure of files. There are lots of ways to do this, and I need not dwell on them. Let's see how to recover from the problem.

First, how can un-erasure be done? Recall that a disk has sectors for data, a FAT, and a directory. The directory points to the FAT, which then maintains a list of sectors connected to the file.

Erasure is done not by blanking out sectors, but instead by changing the first letter of the name of the file, as the name is kept in the directory. DOS doesn't go to the trouble of erasing the sectors, so they just sit there, available for use the next time a new file is created (or an old file is expanded). *As long as you un-erase before a new file is created, the file can be saved.*

The process could be done by hand, with a lot of patience and DEBUG, but fortunately there are products to do this. The best known one is, of course, the Norton Utilities. The original product included UnErase, and as of version 3.1 now includes Quick UnErase (QU). Mace Utilities also include a program like that, called UnDelete.

Simple un-erasures can be done with QU or UnDelete fairly easily. I prefer Mace, as Mace will un-erase nondestructively to another disk. If a number of files have been erased, or new files have been created, then the unerase programs can be fooled. With a program like Mace, you can try the process a few times under different assumptions. With Norton, you cannot do this.

7.8.5. Recovering Accidentally Formatted Hard Disks

A more painful problem. It is more likely under DOS 2.0, 2.1, and 3.0 than in 3.1 or 3.2, but it still happens. As mentioned before, the DOS FORMAT command does not actually erase sector data, but only erases and creates the boot record, the FAT, and the root directory.

So the sectors are all there, but how can they be sewn back together? (A 32 MB hard disk has about 64,000 sectors.) In the case of subdirectories, the problem is solvable. Subdirectory information is maintained in files, and those files are recognizable by their format. Unformatting programs first locate these files, then use the information in them to reassemble data files. The subdirectory files contain a link to the first FAT entry. Since the FAT is a blank FAT, it's only of use to tell us what the first sector was.

Unformatters then read sectors until they come to another file. This means that fragmented files cannot be saved in entirety, and garbage sectors may be tacked on the end of data files. But it's better than nothing, and your chances of success can be improved by running an unfragmenting program now and then (see the next section). It also means that any files kept in the root directory cannot be restored.

Another approach is to keep a compressed copy of the FAT, the directory, and the boot record in a specific physical location. As you keep it up to date, the disk can be restored. Mace uses this approach when it creates a file called BACKUP.M_U. A program, RXBAK.EXE is run periodically and the information is kept up to date. Then, rather than unformatting, just restore BACKUP.M_U and your data is restored entirely.

Another approach is to remove FORMAT from everyone's disk, or create batch files to disallow formats of the hard disk. Sometimes it works, and sometimes it just encourages the curious to experiment. The best defense is a good backup.

7.9. MAINTENANCE: DISK OPTIMIZATION

Following are a few ideas and products to make hard disks a bit more effective and useful.

7.9.1. Fragmentation Of Files

As mentioned before, files can be *fragmented*. This means that a given file may be in sectors 30-40 (one contiguous group) and 101-122 (another contiguous group). Noncontiguous files are bad for two reasons.

First, noncontiguous files take longer to read. The fact that the disk head must go chasing all over the disk to read them slows down disk access. Putting the whole file together makes read operations quicker.

Second, you have a higher probability of recovering an erased file if it is not fragmented. Keeping disks unfragmented just pushes the odds a little further in your favor.

How do files become fragmented? Well, when DOS needs a sector, it basically grabs the first one available. When the disk is new, only a few sectors have been taken, and the rest of the disk is one large free area. New files are all contiguous. But, as files are deleted, they create "holes" in the disk space. As DOS is requested for new sectors it responds by taking the first sector available. This can easily lead to a new file being spread out over several separate areas. You can test this with CHKDSK.

Programs exist which will rearrange the disk so that all files are contiguous and no holes exist. The best known is Disk Optimizer from SoftLogic Solutions. The latest version even allows you to specify most-used files so that they can be moved to the front of the directory, speeding access by a bit. Mace will also do this, with its UnFragment option.

If you use copy protected software, like Lotus 1-2-3, talk to the vendor before trying this. Lotus's annoying copy protection scheme will cause 1-2-3 to be destroyed if you optimize a disk with Lotus installed. If you run copy protected software, de-install it before optimizing, then re-install it.

Better yet — *remove* the copy protection. CopyWrite, from Quaid Ltd., will remove the copy protection from Lotus. I am *not* advocating piracy. But I *am* advocating protecting your investment.

7.9.2. Ever-Growing Subdirectories

Subdirectories never get smaller. The files which contain the subdirectory's directory grows as the number of files grow, but never shrinks if the number of files is lessened.

Directory entries of erased files remain, as mentioned before, unchanged only by a single byte. Holes then exist in directories, which slow down disk access.

"Directory squeeze" programs clean up both of these problems. Mace does it. I'm not sure it really contributes significantly to performance.

7.9.3. Disk Caching Programs

A disk cache works on the principle that a sector used on disk tends to be re-used soon. Given this, why make the disk read the sector? Instead, some memory is set aside, and the most-used sectors are kept there. This can speed disk access to nearly memory speed.

Disk cache programs — Lightning! is probably the best-known, although fairly limited, and I can't recommend it in good conscience — take a while to "get up to speed." Flash! improves on this by giving you a chance to specify files to load straight into a cache. This feature is actually unnecessary: you can load a file into the cache by copying it to NUL:, like COPY BANANA.WK1 NUL:. They don't know when you first boot up which sectors are most likely to be read, so they take a while to "learn" which sectors are most profitably cached. And, of course, if you have an application which jumps all over the disk (unlikely, but possible), the cache will actually slow you down.

What if cached data is sitting in memory and the computer is powered down? Won't that lose data? No — any sector writes are done directly to disk. It's really disk reads that benefit from caching, although a smart cache observes when a sector that the system wants to write doesn't vary from the current disk value. Golden Bow Systems sells a utility called Vcache which seems fairly good, and it is included with Mace. For real performance, however, plan on getting an expanded memory (Lotus-Intel-Microsoft) board. The caching program then uses this expanded memory rather than your important 640K to cache data. Even if you don't run Lotus, these boards are becoming quite reasonable. They can be bought for about \$120 empty of memory chips, and if you're feeling lazy you can buy them stuffed with 2 MB of RAM for about \$300.

As I said before, I would *not* recommend Lightning!. It has two annoying problems:

1. It will not employ AT extended memory (Vcache will).
2. Even if it *does* use Lotus/Intel/Microsoft memory, it *still* takes up 56K of your lower 640K memory. Personal Computer Support Group (the creator) tells me there's no way around it.

Don't buy such a thing if you have a PS/2 or OS/2, however. The PS/2s come with a cache program, IBMCACHE.SYS. (You have to search a bit for it, but it's there. It is actually on a hidden file on the SETUP disk, and you get to it with a batch file. No, I have no idea why IBM did this.) OS/2 has a disk caching capability right in it.

8. FLOPPY DISK DRIVES

Virtually 100% of the PCs or PC clones today have floppy drives. IBM's inclusion of a cassette interface was never taken seriously: when was the last time you saw a program which could be bought on cassette? (I think the Diagnostics was the only one I ever saw.)

8.1. FLOPPIES AND THE PC

Floppies are responsible for a fair amount of what little downtime a PC experiences. They are less troublesome than they were several years ago, not because they are manufactured any better but because we use them less: hard disks are the medium of choice for many home users and most corporate users. The floppy's role nowadays is mainly (1) a program distribution medium and (2) an archival device. As most hard disks can't travel, most computers still rely on floppies to get information from the outside world.

A floppy disk system is a fairly ingenious device in some ways. Part turntable, part cassette recorder/player, they unfortunately rely on a fair amount of moving parts. This makes floppies a weak point in your computer's reliability armor. Fortunately, they are inexpensive enough (\$100 or so) that a fast, relatively cheap repair can be effected by someone with no troubleshooting ability at all.

This chapter will look at the components of a floppy system, and how to test, replace, and adjust these components.

8.2. PIECES OF THE PICTURE

The floppy disk subsystem is, like the rest of the PC, modular. As before, we will use this modularity to allow us to "divide and conquer." The subsystem consists of the drive, a controller board, the cable connecting the drive to the controller, and the floppy diskettes themselves.

8.2.1. The Floppy Diskette

I said above that the floppy drive is part turntable, part cassette. The floppy diskette is the record/cassette. It comes with a built in feature which records lack, however. You may recall, when in your past (or present) audio buff days getting out the Dustbuster and cleaning the record prior to playing it. For some people, ritual like this is comforting. For those people, I must sadly report that no such thing occurs with floppies.

Floppies are stored inside their own Dustbuster: a semi-rigid case lined with fleecy material. The case has a hole cut in it so that the disk can be read/written without having to remove it from the case. **In general, there is no need to clean a floppy diskette.**

A reasonable life expectancy for a floppy diskette is said to be about 3 to 4 years by diskette manufacturers. There are, of course, better and worse diskettes. All diskettes are created in the same location, then tested. The question of single vs double density and single vs double sidedness is determined by the results of the tests. Most diskettes these days are at least double density. Some double density diskettes can format quad density: it's a matter of luck. Similarly, some single sided diskettes can format double sided. Are the manufacturers ripping you off when they charge extra for double sides or higher density? It would seem so: if `FORMAT` likes a medium, it should be all right. Actually, this is not

true. We're not only interested in how well the diskette stores data now, but how well it stores data in the future. Manufacturers have test equipment which can test signal strength of data written to a diskette, and thus can distinguish more than just "data is there" versus "data is not there." You only save literally pennies per diskette like this. Remember that you're using these to save data that cost you a lot of money to generate.

In order to hold the floppy so that it can be spun, the drive "clamps" onto the edge of the hole in the middle of the floppy. Tandon drives ensured good speed control by clamping rather hard, and so with the advent of the PC (the first major computer to use Tandon drives) we slowly saw the growth of floppy drives with hub rings. (Back in the early 80's, companies sold retrofit hub ring kits: hub rings and glue with an applicator.) Nowadays, virtually all double density floppies have them. They are actually unnecessary today, as most drives are half height drives which do not need to exert as much pressure.

8.2.2. The Floppy Drive

The drive itself varies in several ways:

- half height vs. full height
- size (8", 5-1/4", 3-1/2")
- density (double vs quad density)

Half height or full height? Half height vs full height has no effect on data storage: disks read or written with a full height drive can be read or written with a half height. There are even 1/3 height drives — Okidata makes them. Your major concern maintenance-wise is that the half heights are a bit more troublesome to work on, although they seem to need less maintenance. (No data there, just subjective feeling based on machines that I've worked with.) Occasionally you will have trouble putting two half heights in where previously there was one full height. If doing this with a PC or XT (IBM only — the clones have no problem), be sure that you have a mounting bracket to allow you to do this.

Size in the PC world is mainly 5-1/4". Some portables use 3-1/2" drives, and so XTs can be outfitted with them. They can be outfitted with cases so that they fit where a half height 5-1/4" would be. The connection is the same, the controller is the same: you need only inform DOS with the DRIVER.SYS command that the drive is a 720K floppy. The new machines use a 1.44 MB 3-1/2" floppy drive.

Density varies somewhat. Density is measured in *tracks per inch* (TPI). Regular double density drives are 48 TPI drives — the 40 tracks fit in 5/6 of an inch. The 1.2 megabyte floppies used in the AT are quad density drives, 96 tracks per inch, 80 tracks on a disk. 3-1/2" drives are quad density, too.

The quality of the drive also determines how many sectors can be placed on the disk. The table below summarizes these data.

Disk Formats				
Disk Type	Capacity	Tracks	Sectors/Track	Bytes/Sector
360K DSDD Floppy	360/320K	40	8 or 9	512
1.2 MB Floppy	1.2 MB	80	15	512
3-1/2"	720K	80	9	512
HD 3-1/2"	1440K	80	18	512

Disk drives, when installed, must be *terminated* properly. What this means is that the A: drive must be installed with a special resistor (it looks like a chip, in a DIP package). If the drive isn't terminated, the controller won't talk to the drives properly.

Many things can go wrong with disk drives. You can learn to fix them (we'll talk about it in this chapter), but they are sufficiently cheap that you could reasonably consider them disposable. In any case, be sure to have a few on hand for spares. Fixing drives takes time, and you may not have that time when a drive dies.

8.2.3. The Disk Controller Board

As with other devices, the floppy drive needs a controller board. On the AT, this is also the hard disk controller. On the PC, you can buy combination hard/floppy disk controllers for about \$150. A floppy-only controller for the PC is about \$30. Again, controllers can be fixed, if you've got time and you're patient (most chips on controllers are *not* socketed), but it can't hurt to have a spare or two around for quick diagnosis.

You can swap the controller board in just a few minutes. There are generally no DIP switches to set. You'll see a ribbon cable extending from the edge of the controller card. This ribbon connects to the disk drives. Disconnect it from the controller card. It *should* be keyed so that you cannot connect it upside down, but check at this point. If there is any ambiguity, draw a picture. Take a magic marker and write "up" and an arrow on the connector. Do whatever makes re-assembly easy. Then unscrew the mounting screw and remove the disk controller board (you have, of course, turned the power off before doing any of this). Install the spare in the reverse manner and try to reboot. If the problem goes away, fix the controller or throw it away and buy another one. This is an effective diagnostic procedure that literally only takes minutes.

8.2.4. The Cable

The drive is connected to the controller by a 34 wire ribbon cable. The cable is usually "keyed" between lines 4/5 and 6/7 so that it is inserted correctly. Most of the cables have three edge connectors: one for the drive controller, one for drive A:, and one for drive B:.

The connector from the drive controller to B: is a "parallel" cable — i.e., pin 1 on the drive controller is connected to pin 1 on the floppy, 2 connected to 2, etc. The connector to A: has a twist in it, however. It is this twist that identifies A: from B:. Here are the unusual connections on the A: side:

Pin on Controller, B:

10
11
12
13
14
15
16

Pin on A:

16
15
14
13
12
11
10

Use this information to either make up new cables or test existing cables for continuity. The ribbon cable has three connectors: one on one end and two on the other end. The lone end goes on the controller. The one in the middle goes on drive B:. The one on the end goes on drive A:. This is why A: must be terminated — the furthest drive from the controller must be terminated. If you're likely to forget (as I am), get a magic marker and write "controller," "B:," and "A:" on the appropriate connectors.

Testing a cable: before you get out your VOM, however, do it the easy way. Keep an extra cable. The cables are about \$20 - \$30 from any mail order house or computer discounter: keep an extra one around. Swapping the cable is easy. Testing continuity is a pain.

8.3. MAINTENANCE

For all of their moving parts, disks behave rather well. I still have the same A: drive on my original PC as when I bought it in August 1981. A few simple procedures will greatly extend the life of your drives and diskettes.

8.3.1. How Often Should You Clean The Heads?

Like audio cassette drives and videotape recorders, floppies have a magnetic read/write head which collects magnetic oxide. It is a good idea to remove this buildup now and then. This has led to floppy head cleaning kits.

Like all good PC owners, you have probably purchased a floppy head cleaner to ward off floppy evil spirits. Once procured, however, the first question is: how often should I clean the floppy heads?

The head cleaner people say to do it weekly. Some people cynically claim that this is because they want you to use it up and buy some more. Other people claim that cleaning this often would be disastrous, as the procedure wears away a little of the head each time — you always have a clean head, but it lives a shorter time as a result of the cleaning process. The books are split on the question.

Personally, I clean my heads every 12 months whether they need it or not. I have some computers which I've *never* cleaned the heads on, and they work fine. My recommendation: clean the heads once per year, and whenever you start experiencing read/write errors.

Remember another head-saving tip: on full height drives, keep the doors open when you are not using the drive. Otherwise, the two heads grate against each other. Some people think it's smart to leave the drive doors open, as an anti-dust measure. The doors are *not* dustproof by a mile, so it's a silly idea. Leave them open and save your disk heads.

8.3.2. Environmental Factors To Protect Disks And Diskettes

We've all seen the "do not" cartoons on the back of the diskette jackets. They lead you to believe that floppies are very fragile items. As those of us in the real world know, this isn't really true. Yes, they must be taken care of, but you needn't get crazy about it. Don't put them on the radiator, or leave them on a shelf that gets three hours of direct sunlight every day. Don't store them under the roof leak. Given the choice, store them upright stacked left to right rather than on top of each other.

Remember thermal shock. If your portable computer has been sitting in the back of the car in freezing temperatures overnight, bring it in and let it warm up before using it. Just a little heat expansion/contraction can temporarily realign your drives. More extreme temperatures can damage the diskettes: vendors claim that diskettes should never be stored below 50 degrees nor above 125 degrees. (This is probably paranoia: I order software through the mails regularly, and when it is delivered in January it often sits outside in freezing temperatures for the entire day. I've never had a problem. When it sits out in the rain, or when the postal carrier bends it, that's another story.)

Dust, smoke, and dirt can cause damage to the head and/or to a diskette. As before, don't smoke around the machines. It's nice to work with the windows open on a sunny Spring day, but there's a lot of pollen and road dust in the air. To see this, clean your windowsill thoroughly, then leave the window open for a day. Check it the next day. On a normal Spring day when the weather has been dry, you should see a noticeable film after just one day. You sure don't want that on your heads.

8.3.3. A Disk Drive Tester

Part of maintenance is monitoring. Several disk drive test programs are on the market today. Verbatim Corp's Disk Drive Analyzer is probably the most rudimentary, but it is cheap and easy to run. You can run it now and then to be alerted to disk speed or alignment problems. For \$40, it is a very rough diagnostic tool. Others can be bought for much more which offer much more information.

8.4. PROBLEMS AND DIAGNOSIS

In a minute, I'll talk about repairing components of the floppy system. But first — how do you know what's wrong? Let's look at some common symptoms and possible causes.

Suppose the drive is malfunctioning, and one of the following things are occurring:

- refuses to read
- refuses to write
- reads, but grinds a lot and displays "retry?" messages
- you see a "601" message, then the computer goes to Cassette BASIC

Possibilities fall into two categories: software and hardware.

8.4.1. Software Possibilities

- Bad applications program
- File destroyed by software bug or incorrect power-down
- FAT damaged
- Directory damaged
- Boot record damaged
- File accidentally erased

These are discussed in the chapter "Hard Disk Maintenance and Failure Recovery." The same software that will help on hard disks will generally work on floppy drives also.

8.4.2. Hardware Possibilities

Assuming that the drive has worked before, the following are possible causes. When I say, "remove the drive," and you say, "*How?*" — be patient. I'll cover all that by the end of the chapter.

8.4.3. Drive Head Is Dirty

Try cleaning the drive head. This may solve the problem.

8.4.4. Diskette Is Bad

It is fairly common for diskettes, particularly heavily used diskettes, to develop bad spots. If you always use the same diskette as your work diskette, try another one. If the problem goes away, use a disk tester like Norton's DiskTest or the Mace Utilities DIAGNOSE program on the original disk. Recover what data remains and discard the diskette. (You *do* keep backups, don't you?)

8.4.5. Cable Is Loose

Sounds silly, but "is it plugged in?" Expansion/contraction and vibration can loosen the power and data cables on floppies. Maybe they weren't installed snugly to begin with.

8.4.6. Cable Is Bad

In truth, I've never gotten a bad *floppy* cable. But I've seen bad cables of other types. If the floppy worked yesterday, the cable probably works today. But it only takes a two minute swap with another cable to find out for sure.

8.4.7. Controller Is Bad

Generally, you will see a "601" error on powerup if the controller is malfunctioning. You could run the IBM diagnostics to test the controller, but, again, in some cases just swapping the controller is a quicker way, as was mentioned earlier.

8.4.8. Power Supply To Floppy Malfunctioning

At the back of the floppy, where the line from the power supply connects, you will see four solder points. They will probably be labelled, left to right, "+12V GND GND +5V". These are the power lines that the floppy needs to operate. You can test these while the floppy is operating with a VOM. Put the black probe on "GND" and the red probe on either "+12V" or "+5V". You should see power in the range [4.8, 5.2] for the 5 volt pin, and [11.5, 12.6] for the 12 volt pin. If the values are outside of this range, the power supply is faulty.

While we're on the subject, I should mention that some "disk failures" are actually the fault of power spikes. Power protection is a good move here, as always.

8.4.9. Drive Is Malfunctioning

The floppy drive itself can malfunction for a variety of reasons:

- head is dirty
- head is mis-aligned
- rotation speed is wrong
- drive door is broken or cracked
- on-drive electronics have failed
- track zero adjustment is wrong

At this point, I'd do the following things:

1. Reboot. The disk parameter table may be bad.
2. Check the diskette. Try another floppy.
3. Clean the heads on the disk.
4. Swap the drives. Check that connections are solid.
5. Try a new A: drive.
6. Swap the controller.

You don't want to spend forever figuring out which is the reason. You just want (at this stage) to determine if it is *something* on the drive which is causing the problem.

The simplest approach to this is to simply swap the drive. For details, see the later sections. The steps are as follows:

1. Remove drives A: and B:
2. Remove the terminator chip from A: and put it on drive B:
3. Re-install the drives, putting A: where B: was and vice versa

If you can now boot from drive A:, and read/write data, then the former A: is at fault. Replace or repair.

8.5. COMMON REPAIR POINTS

Once you have a notion of what part of the floppy subsystem has failed, how do you replace or repair it? This section addresses that.

8.5.1. Is It Worth Repairing Floppy Drives?

As before in this text, I want to open this section by saying that if you only learn how to replace the controller board, cable, and drives, and nothing about repairing the components themselves, then you're doing fine. Floppy drives are very cheap — about \$100 — and can be replaced without repair. Even if you want to repair the drives, there are firms that will fix floppy drives for a flat fee of \$35-\$45. This may be more cost effective than spending your time trying to repair it.

The minimum that you should be able to do is:

- configure and replace a floppy drive
- replace a floppy controller board
- replace the drive cable

Extra capabilities could be:

- replace a broken drive door on a full height drive
- adjust speed on the drive
- repair and recover a dirty or damaged floppy

8.5.2. Removing Floppy Drives

Floppy drives are removed in three steps:

1. remove power connection
2. remove data connection
3. remove screws from mounting brackets

The floppy drive is secured to the computer case with two mounting screws. These can be a pain to remove, as the two screws on the left side (the A: drive) are usually close to expansion boards. To loosen these screws you sometimes need to remove all of the circuit boards. A useful tool here is an *offset screwdriver*. This is a screwdriver with a bent handle. Using it can keep you from having to remove all of the expansion boards.

On the back of the floppy is a connection from the power supply. It has four wires attached to a connector: a yellow, two blacks, and a red. The connector is keyed so you cannot connect it backwards. There are two of these (four on some power supplies): one for each floppy drive. Many power supplies label them "P10" and "P11," P10 on the A: drive and P11 on the B: drive. Actually, it doesn't matter which is connected to which. It's the twist in the drive cable which decides which is the A: drive. The connector may be under the back of the printed circuit board. Disconnect it with a gentle pull. You may find it easier to pull the floppy out just a bit so you have room to reach around back and disconnect cables.

A blue or gray edge connector on the data cable connects to the back of the drive. Remove it.

Now the drive will come out fine. As you remove the drive, be careful of parts on the drive that stick up. Don't break them off as you remove it.

8.5.3. Installing A New Floppy

Installing a floppy is just the reverse of removing one, except the drive must be configured. To configure a new floppy, you must:

- terminate the floppy properly
- set the drive select jumper

The last drive on the drive cable — A:, in the case of the PC — must be *terminated*. Termination just means that a resistance must be present across two electronic lines or the drive controller thinks that there's nobody home. You're only supposed to terminate the last drive, but I've terminated both drives and everything's turned out fine. You must be sure to terminate at least drive A:. Termination is done with a terminating resistor. It looks like a chip — in fact, it is a resistor in IC's clothing — and comes on drives socketed. To un-terminate, just remove the chip. Remember: **Drive A: must be terminated.**

Second, the drive must be selected. This means that certain jumpers must be added or removed. On full height drives, there is an area that looks like a socket for a chip. "Pin 2" and the pin across from it must be jumpered. A regular staple works fine, and is probably the jumper of choice for 90% of the PCs installed today. For other drives, look for "DS1" (Drive Select 1) or a marking like that. On Shugart half-heights, use DS2. On other drives, it's DS0 through DS3: on Shugarts, they're labelled DS1 through DS4.

Sometimes you will see two jumper positions "RDY" and "CD". In general, a jumper should be on one or the other, but not both. RDY is used on a PC or XT. CD is used on an AT.

8.5.4. More On Drive Select

Only read this if you are curious about the twist in the cable, and why DS1 is selected.

Back in the old days, daisy chained drives spoke, one at a time, over a shared cable. They knew what address they had by their Drive Select jumper. The first drive was DS0, then DS1, etc.

IBM decided (perhaps rightly) that actually requiring users to select DS0 for drive A: and DS1 for drive B: was a bit taxing. Was there an easier way? As it turns out, by putting a twist in lines 10 to 16, *both drives can be DS1, but the one with the twist looks like DS0*. The twist, then, makes life easier for the PC installer. While I have never tried it, this suggests that two more floppies could be easily added with a non-twist cable and explicit drive select jumpers for DS2 and DS3. **Warning:** that last line is mere speculation. I don't endorse it.

All of the above is from people who claim to know. It seems to tell the story, but I wonder why I can't use a PC with a no-twist cable and an A: drive selected for DS0 (or DS1 or DS2 or DS3, for that matter — I've tried them all). No one seems to know.

8.5.5. Replacing The Drive Door (Full Height)

When the PC first came out, people soon learned that the plastic doors on full height drives tend to break easily. The spring-loaded doors snap open, sometimes destroying themselves in the process. If the door can't be closed, then the drive won't even try to read. Sometimes the door will only crack, leading to the mysterious case where the drive door is closed, but DOS says the drive is not ready.

Replacing the drive door is fairly easy. Finding a replacement is hard. You can get replacements from:

Workman Associates
1925 East Mountain Street
Pasadena, CA 91104
(818) 791-7979

First remove the screws that hold the drive to the computer case. Then pull the drive out a bit.

The drive is held to the spindle by a metal bar. The metal bar is held to the door by two screws, right behind the drive face bezel. The bar is spring loaded, so hold it down with one finger while removing the old door. Then put the new door in, and fasten to the bar with the two screws.

8.5.6. The Floppy Drive Electronics

On each floppy of whatever make, there is at least one printed circuit board. As always, these may be defective, but they tend to live fairly long if they survive the first few weeks. Unfortunately, there is no way to "burn in" these boards without burning out the drive stepper motor.

They are not difficult to change, but I'd not recommend it. I have never seen a floppy with a bad PC board. Further, it is very difficult to get replacements. Some floppy drive manufacturers will not sell you replacement doors, boards, etc. unless you are a repair facility licensed by them. Licensing usually requires some franchise money and attendance at their repair seminars. As I say, for \$100 throw away the floppy. Or replace it and keep the bad one for parts.

8.5.7. Saving A Dirty Floppy

Earlier, I said, "in general, there is no need to clean a floppy diskette." That's "in general."

I *have* seen some occasions when cleaning a floppy is appropriate. For instance, after flood damage, or when the odd glob of chocolate Haagen-Daaz gets onto the floppy. The instances when you clean a floppy are extremely rare.

Here's how to do it. Cut across the top of the floppy jacket. Carefully remove the floppy and place it on a clean, soft surface — a paper towel will do, but cheesecloth is preferable. The floppy can be wiped clean with a damp soft rag with warm-to-hot soapy water, then *gently* wiped dry with a soft clean cloth, and dry on a flat surface. I once saw an article claiming that the floppies could be dried by stringing them up through the center hole, but I doubt that this would work well: the floppy could dry in a bent shape.

Wait until the floppy is dry. When I say *dry*, I mean it. A damp floppy could flake wet oxide bits onto your drive head and destroy it or seriously incapacitate it.

Then take a blank floppy (this is the "sacrificial" floppy) and remove it from its case: throw it away. Use its case for the cleaned floppy. If you make a clean cut with a razor across the top of the disk jacket, then the jacket can be re-used.

This, by the way, is also useful sometimes when a floppy is mailed and the mailer is bent. (Some postal employees take "do not bend" written on the outside of a diskette mailer in the same way that kids take "do not touch.") If you can't unbend the case by running it along the table edge, take the floppy apart and steal another case.

If you've been successful here, *copy the floppy immediately*. If it is a distribution disk, send away for a replacement just in case. This is only a disaster recovery procedure — not a maintenance procedure.

8.6. ADJUSTING FLOPPIES

As they are mechanical devices, floppies occasionally need adjustment. Besides head cleaning, floppies can have their speeds adjusted and heads aligned. Head alignment is a bit tough to do without investing in an oscilloscope, an expensive piece of equipment, but a rough alignment can be done without it.

8.6.1. Disk Rotation Speed

Floppies rotate at 300 rpm. The 1.2 MB floppy drives on the AT rotate at 360 rpm. Small variations in this speed (plus or minus five percent) are okay, but greater differences can make the floppy nonfunctional or dangerous.

Dangerous? Yes. A few years ago, I was asked to install a copy of dBase III version 1.0 on an XT. It was a perfectly legal copy. When I went to install it, it informed me that it was an illegal copy, and refused to load. I tried installing it on other machines. It didn't work. Ashton-Tate first told me that I was lying or mistaken, and finally (after threats of violence) sent me another master disk⁸. It took a while to reconstruct the problem.

What had happened was this: the floppy on the XT had a drive that was slightly out of correct speed. Not so different as to render it unusable for normal uses, but enough to upset the finicky copy protection system. The particular system being used at the time — Prolok, by Vault Corporation — would *permanently alter a disk which it once saw as an illegal copy*. Once Prolok thought a disk was bad, it would alter the disk so that it would *always* seem bad. That was why the disk wouldn't run on other XTs. (How did I eventually prove this? I have some friends at a software house using Prolok. They Prolok-ed some disks, and we tried to load them on the XT. Same problem.) Copy protection is almost a quaint remembrance, but not yet. It pays to be able to test and adjust disk speed.

8.6.2. Testing And Adjusting Disk Speed

First, understand that you will probably never have this problem if you buy half height drives. They are direct drive, rather than belt driven. Some of the newest ones don't even have an adjustment switch — logic on the drive automatically adjusts the speed. No maintenance.

There are two basic ways to test disk speed: a speed tester program and stroboscopically.

Speed tester programs exist for 5-1/4" floppies. The two most popular are Verbatim's Disk Drive Analyzer and Readiscope. They indicate either graphically or numerically whether the drive is too slow or too fast. You remove the drive from its bracket, then put it on its side on the top of the power supply — a convenient flat surface. Reconnect the cables and run the program, adjust the speed and run it again. Keep doing it until you get the speed just as you like it.

There is a public domain floppy RPM tester program on your program disk called SPEED.EXE.

Speed on a floppy is adjusted by a variable resistor (varistor) on the floppy. It is the "tweaker" screw with the red epoxy on it. Just break off the epoxy and adjust with a screwdriver, then put a dab of nail polish on it to hold the setting.

The varistor is on a small printed circuit board sitting upright on the back of a Tandon floppy drive. The CDC drive has it on the circuit board on top. A Shugart half height has it on the top circuit board on the right towards the front. The speed adjustment varistor is usually easy to find, as it is often the *only* varistor on the floppy.

The other approach doesn't involve software. The bottom of *some* (not all) disk drives have a pattern which looks like a test pattern. It is two concentric sets of spokes, one inside the other. These are used for a stroboscopic speed test. Here's how it works.

First, remove the drive. Then get the motor running: either issue DOS commands or see below. Put a stroboscope flashing at 60 cycles per second (hz) on the spokes. The outer set of spokes will look like they are standing still when the speed is correct. (The inner set is for 50 cycles per second.) The easiest and cheapest 60 hz strobe is a common fluorescent bulb.

⁸Thankfully, A-T has now commendably removed copy protection from their products.

How do you make the drive motor turn continuously without software commands? Check the manufacturer for specific drives (I hope you've inquired about a technical manual). Here's how for a Tandon or a CDC. On the Tandon, run a jumper from TP10 to TP13. TP10 is a pin sticking up from the circuit board on top of the drive, near the left front of the drive. TP13 is on that board near the left rear, near the power connector. On the CDC, you will see a long strip of pins on the rear of the top circuit board. The sixteenth from the left (assuming that you are looking from the front of the drive back) should be jumpered to ground.

Remember to secure the varistor when you're done adjusting it. Use electrician's epoxy, or just regular nail polish will do it. (It also comes in a variety of colors, unlike boring red-only epoxy.)

8.6.3. Head Alignment

You've probably heard of floppy disk head alignment. A fair number of disk troubles get blamed on this. Many shops will align drives for about \$35, not an unreasonable price considering that it can take an hour or more with the right tools.

Alignment is not simply setting one single screw. It is actually several adjustments:

- Head Radial Alignment
- Track Zero Sensor Adjustment
- Carriage Stop Adjustment
- Index Timing Adjustment (skew)

Picture several tracks of data on a diskette. They are concentric rotating circles. Now imagine the head. It wishes to straddle one circle at a time, reading and writing data as the track goes around. If it gets lost and tries to work between the circular tracks, its *radial alignment* is off.

As the head straddles a track, it meets oncoming data head-on. If it is not angled exactly forward, so it faces the circle at a tangent, then its *azimuth* is correct. The head can only be about a quarter of a degree off.

The head moves in toward the center and out toward track 0 along a straight line. If this line does not run directly through the center of the circles, the disk has a bad *skew*, called *index-to-data-timing*.

With an oscilloscope and a special alignment disk, all of these adjustments can be done. As this is pitched at a non-technical user, I don't propose to show you how to use an oscilloscope. Instead, we'll discuss alignment with a software program.

8.6.4. Aligning The Heads With A Test Program

Strictly speaking, you can't align heads with a software program. The reason is simple: there are just too many things that can be mis-aligned. All a software program can tell you is whether you're aligned or not. If you're not, you just adjust something else and try again.

If you'd like to try, however, Readiscope offers directions on adjusting radial head alignment. It draws a little picture on the screen representing your drive's head and the track that it is trying to read. You adjust the head and the picture moves.

This is a nice program, and it does what it does well. The problem is that it is no more likely that you will have a bad radial alignment than an azimuth or skew problem, and the program kind of leaves you hanging on that point.

9. PRINTERS AND PRINTER INTERFACES

Printers can be a real maintenance headache. As they produce tangible results (pieces of paper), malfunctions with printers can be more upsetting than their wholly-electronic brethren. For example, try telling Lotus that you've got an Okidata printer when you have an Epson, then print a graph. The printer will start spewing out pages, as if angry for being misrepresented.

Most of *my* printer problems — the vast majority — are software-related issues. The hardware problems seem more often to be cable or printer interface problems. I use mainly dot matrix and laser printers. Daisy wheel printers *are* prone to hardware failure, and I frankly recommend that you do not buy them. As of early 1987, a 500 cps laser printer could be bought for about \$1200, the price of a 50 cps daisy wheel printer with a sheet feeder. And it's a whole lot quieter.

9.1. COMPONENTS

As always, a controller (the parallel or serial interface board), a cable, and a printer. The parallel port is simple enough that it is commonly included on other expansion boards.

9.2. MAINTENANCE

A few things can be done to maintain printers of all types. Vacuum out the paper chaff periodically from the inside of the printer. Determine if there is a belt tightening mechanism for the printer — usually a motor moves the print head via a belt. Find the correct tension values. Keep a replacement belt on hand. (Believe me, they're no picnic to find in a hurry).

With a dry, soft cloth, clean both the paper path and the ribbon path. Most manufacturers suggest cleaning every 6 months. The ribbon path can build up a film of inky glop which causes the ribbon to jam. To do this, go to a drug store and buy a dispenser box of 100 clear plastic gloves. Use them when working on the printer (but not chips and boards — that plastic can build up some mean static) so that you don't have to wash your hands for hours to remove the ink.

Most printers do not need to be lubricated in everyday use. In fact, oil can do considerable damage if applied to the wrong places. If you thoroughly disassemble the printer, then you will probably have to lubricate various points as you re-assemble. If you intend to do this, I strongly recommend that you get a maintenance manual from the manufacturer.

9.2.1. Daisy Wheel Printers

Daisy wheel printers require that the daisy wheels be replaced about annually. Some variations, like Spinwriter thimbles, last longer, but must be replaced eventually. (The Spinwriter also needs a striker shield replaced every few months. It is a thin flat piece of plastic which sits near where the thimble strikes the paper.) Their belts must be replaced now and then. The gap between the printing element and the paper must be aligned regularly — a little too regularly, actually. (Get a laser.)

9.2.2. Dot Matrix Printers

The expensive part of a DM printer which dies is the print head. Print for too long a time and the head just burns up. This is not as much of a problem for the newer printers as it was for the old Epsoms and Okis. The models out today have a thermistor which shuts down the printer temporarily if the print head overheats. If the thermistor becomes ill, the printer shuts down regularly. Generally thermistors are pretty robust, but if such a thing happens try changing the thermistor first: it's a lot cheaper than a print head.

On older Epson printers I am told that replacing print heads is not really a reasonable fix, as the print heads usually take the circuit boards with them when they burn up.

Replacing print heads is not economical on many printers because of the high price that manufacturers charge for replacements: often it is a significant percentage of the entire cost of the printer. My Okidata 92, which I purchased years ago for \$370, would need a \$125 print head if I needed to replace it.

Dot matrix printers are, in general, very reliable. But keeping the print head cool is vital. Don't stack up things on or around the printer.

A variation on the typical dot matrix printer is the ink jet printer. Rather than hammering at a ribbon, the ink jet squirts a narrow jet of ink at the paper. This is very quiet, but the jets tend to clog, leaving partial letters on the page. The answer here is simple: remove the cartridge and push on the ink sack with a long thin tool, like a straightened-out paper clip. The ink will push out the small holes, unclogging them. With the HP ThinkJet, it's a pretty regular procedure. The Seiko color ink jet has the same problem.

9.2.3. Laser Printers

The laser printer is very similar to a copy machine. Having said that, it's amazing that they are as reliable as they are.

The most common laser engine is made by Canon. The HP LaserJet, the AppleWriter, the Canon A1/A2, the QMS Kiss, and others are all built around the Canon engine. These need no maintenance except for a new cartridge every 5,000 copies or so. The cartridges cost in the neighborhood of \$70, and, according to HP, contain all that is needed for a routine maintenance. So, you are performing a routine maintenance every time that you change your cartridge.

Lasers require proper ventilation, and a fair amount of power. Other than that, don't pour any Cokes in them and they last a long time.

An important thing to understand about laser printers is that they are *page printers*. This means that they won't do anything until the page is full. You can force a laser to eject a page with the following batch file.

Create the batch file. Bold-faced is what the computer types, regular type is what you respond.

```
C>copy con: eject.bat
echo L> PRN
<F6>
1 file(s) copied
C>
```

Where you see "L", generate this by holding down the Ctrl key and typing "L". Where you see "<F6>", just press the F6 function key.

From this point on, you can just type EJECT from DOS and your printer will eject the current page.

9.3. COMMON PROBLEMS AND TROUBLESHOOTING

9.3.1. Isolate The Problem

As always, try to isolate the problem: something in the computer or its software? The printer interface? The cable? The printer? Is the printer plugged in, cabled, and *on-line*?

The steps I use are:

1. Check if it is on-line, plugged in, has paper, turned on.
2. Cycle power switch on printer, reboot, and retry.
3. Check that the software is configured for the printer.
4. Swap cable to test it.
5. Swap printer to test it. (Use the same type of printer.)

The first thing I'd try would be to cycle the power switch and restart the software. One time I was experimenting with graphics on a daisy wheel printer, setting the vertical and horizontal motion increment to a very small value, then using periods (".") for graphic points. (I got some impressive results, for a dumb daisy wheel.) Anyway, I forgot to reset the motion increments. A secretary came by about an hour later, saying, "Mark, were you messing with the printer?" She held out a letter which the printer had typed in a space about 1 inch by 2 inches.

- Does the printer have paper? One of my favorite tricks is to check that the main paper tray is full, then accidentally and unwittingly select the empty alternate paper tray.
- Run a self-test on the printer.
- Swap the cable.
- Try another printer on the interface.
- Do a DIR from DOS, then try a Printscreen.
- There's lots of buggy printer-driving software: have you ever seen this software work on this printer?

Check the ribbon on the printer. Is it worn out in one spot, near the head? Does the ribbon move — check the ribbon transport mechanism. Make sure it isn't jammed. Check the paper clamp — if you are using tractor feed, it should be disengaged. Otherwise, it should be engaged.

9.3.2. Cable Lengths

The role of cable lengths in noise and interference has been discussed before in this text. But another problem is overly long cables. Serial and parallel cables aren't supposed to be made up longer than 50 feet, and it's probably not a good idea to make up a parallel cable longer than 6 feet. If you're using long cables and getting mysterious errors, the cables may be the culprits.

The cable won't refuse to work — you'll get odd transient errors.

How to solve this? One solution is offered by Intellicom. Called the Long Link Parallel Interface, it extends parallel cables by up to 7,000 feet. It's \$199 from Priority 1 in southern California.

9.3.3. Setup DIP Switch Problems

Generally, printers must be configured. The most common configuration problem is the "AUTO LF," or automatic line feed. This says, "every time you get a carriage return, assume that there's a line feed with it." If your computer sends line feeds anyway, everything comes out double spaced. This is generally adjustable with a DIP switch.

Another configuration option — a more and more common one — is emulation mode. Many printers nowadays will emulate a Diablo 630 printer, or an Epson MX80, or an IBM Graphics Printer. If you've got your Hewlett Packard ThinkJet printer set up for IBM emulation, don't tell your software you've got a ThinkJet — tell it you've got an IBM. This sounds simple, but you'd be surprised at the number of people that get tripped up on that one.

This final one is an indication of how international the electronics business is. Many printers speak foreign languages. If you set up your printer for British, you may get the Pounds Sterling sign rather than a dollar sign.

9.3.4. Port Problems

As you know, printers can have either a serial or parallel interface. On the PC, serial and parallel interfaces differ only by male connectors for serial interfaces and female connectors for parallel interfaces.

They are radically different, however. The parallel interface uses different voltages and handshakes than the serial interface. Printers can be bought with either serial or parallel interfaces. Given the choice, take parallel. It's a cleaner and faster interface. As laser printers get faster and, more particularly, support higher resolution, more high speed interfaces will appear. For example, Apple uses an AppleTalk 230,000 bps interface to an Apple LaserWriter.

Remember to use the screws on the connector to secure your cable. I saw a situation wherein an Okidata printer printed consistently incorrect characters. I tried to understand the problem by comparing the ASCII codes of the desired characters to the codes actually printed. I found in each case that bit 6 was *always 1*. It turned out that the wire for line 6 was not fully seated. Securing the connector did the job.

I found a similar problem with a broken wire in a cable. Here's an example. Suppose I try to print "Hello" but get "Iekko". Compare the codes of the desired and actual characters:

Desired Character	Code	Actual Character	Code
H	01001000	I	01001001
e	01100101	e	01100101
l	01101100	k	01101101
l	01101100	k	01101101
o	01101111	o	01101111

The "e" and "o" aren't affected, but "H" and "l" are. Notice that in all cases the low bit is "1".

If you are using a serial interface, are the communications parameters set correctly? There are four:

- Speed (1200, 2400, 4800, 9600)
- Parity (Even, Odd, or None)
- Number of data bits (7 or 8)
- Number of stop bits (usually 1 or 2)

You'll find the parameters in the technical manual of the printer. Then construct the DOS commands

```
MODE COM1: speed,parity,data bits,stop bits,p  
MODE LPT1:=COM1:
```

In the first case, an example would be

```
MODE COM1:9600,N,8,1,P
```

meaning "9600 bits per second, no parity, 8 data bits, 1 stop bit." The "P" means that it's a printer.

If you're having port problems, did you install something recently? Could something be conflicting with the printer port? Recall the story of the 5251 emulation board. You'd never imagine it, but the terminal emulation board was killing the port. Is there another printer port? Are they both set for LPT1:?

I know that I've said in the past, "don't bother fixing boards," but you might give it a try with a printer port. The chips in a printer port are usually fairly robust SSI or MSI chips, difficult to damage with static and solder. If I were looking for some experience repairing circuit boards, I'd start from printer ports. There's no large, delicate, expensive chips. And if you blow it, big deal — you were going to get a new port anyway. This does not apply if the port is on a larger, multifunction board: you could damage other chips while attempting to fix the parallel port.

9.3.5. Software Problems

Is the program configured for the printer? If you just replaced your old Qume daisy wheel with a HP LaserJet, then the software won't work unless you tell it that you have a LaserJet.

Know your printer. If you're the local PC guru, get to know the bizarre escape sequences that put the printer through its paces.

On the DOS disk, there is a program called GRAPHICS.COM. It allows you to use the `PrtSc` key even when a graphic image is on the screen. (The original PC, without GRAPHICS, will simply ignore graphic data.) Don't install GRAPHICS unless you have an Epson or IBM printer. It won't work. This is because graphic printing commands vary widely among printer types. If you own a HP LaserJet, an Okidata, or some other non-IBM printer you will need replacement software. Contact the manufacturer, or learn some assembler and write such a routine yourself. I've had to do it for two printers, and it's not impossible. It *does* take some time the first time around, however.

9.3.6. The Mysterious Timeout

Sometime the computer will sense that the printer is ignoring it: the printer will "time out." When the computer says, "Abort, Retry, Ignore?" you say "Retry," and it works fine. How can you address this problem? Simply add the following DOS command to your AUTOEXEC.BAT:

```
MODE LPT1:.,P
```

This instructs the computer to retry forever. This means, of course, that you must be sure to have a printer connected, or the first attempt to print will lock up the computer.

Beware, however. Remember the Compaq/`PrtSc` problem in an earlier chapter? If you do not have a printer attached, and use `MODE LPT1:.,P`, the computer will wait forever to print. The three finger salute or the Big Red Switch are your only alternatives.

9.3.7. The Weather

Everyone talks about it, but...

A printer repairman told me about a day he'd had the previous October. He said that all over town a particular model of printer was failing left and right. He couldn't figure it out. We thought about it. Around the middle of October, we turn on the heat in Washington. That dries out the air and, in turn, the items in the work area. Chips don't mind being dried, but what about capacitors? Could a paper-type capacitor be malfunctioning because it was drying beyond a certain point?

A repair memo came around from the manufacturer a couple of months later. Sure enough, a particular capacitor didn't like it too dry. The answer: either put a humidifier near the printer, or change the capacitor to a similar, less dry-sensitive replacement. Moral: be suspicious when the seasons change.

10. MODEMS AND SERIAL INTERFACES

10.1. INTRODUCTION

When computers first appeared on the scene, each was an island. Computer to computer communication mainly came in the form of moving tapes or card decks from one machine to another. We now jocularly refer to this method as *sneaker net* or *AdidasLAN*.

Many communications options now exist, but they fall into three broad categories:

- Asynchronous Communication
- Synchronous Communication
- Local Area Networks (LANs)

Communications troubleshooting is a separate course all in itself, so I won't try to cover the entire topic. I'll introduce the basic components, then discuss some very simple diagnostic techniques.

10.1.1. Asynchronous Communication

Async communication tends to be:

- low speed (generally 2400 bps or under)
- inexpensive
- minimal error checking
- over phone lines

Common applications are computer remote bulletin board systems (RBBSeS), PC to PC transfer over long distance, and communications services like Dow Jones and CompuServe. Equipment needed are generally:

- a modem
- an RS-232C interface on the computer (serial port)
- an asynchronous communications package like Crosstalk or Smartcom II
- a cable

10.1.2. Synchronous Communication

Synchronous communication tends to be:

- micro to mainframe communications
- low to medium speed (2400 to 19,200 bps)

- relatively expensive
- good error checking
- phone or local connection

Synchronous communication is seen in mainframe shops, particularly IBM. Terminal emulation boards like the IRMA products are for synchronous communication. We will not discuss these here. A good approach to troubleshooting is to ensure that the micro-to-mainframe product that you buy fits into the diagnostic and control architecture of your mainframe network. Then the system which is currently in place can track your PC *as a terminal* on the network, just like any other communications device.

10.1.3. Local Area Networks (LANs)

LANs are:

- very high speed (1 Mbps+)
- used mainly for resource sharing
- expensive (\$2000/PC and up)

LANs tend to be tied to a single vendor. Only with the oldest and most established can you buy parts from various vendors and put them together into a single system. As I said before, that is a program in itself, and will not be discussed here.

10.2. COMPONENTS

10.2.1. Asynchronous Port

On the computer end, the computer must be able to speak the language of asynchronous communications. A device to allow this to happen is called an asynch port, asynch adapter, communications port, or RS-232C port. (RS-232 and RS-232C are the same thing.)

The connector type used is generally the standard 25 pin DB25 type connector, such as you would find on the back of most modems. Since the advent of the AT, some computers are appearing with 9 pin versions. You won't miss the other 16 pins — asynchronous communication doesn't use them anyway. You *may* have a cabling problem, however. We'll discuss cables later.

Most folks don't buy an asynchronous adapter on a board by itself. Instead, it usually appears on a multifunction board like the Quadram Quadboard or the AST SixPakPlus.

10.2.2. The Cable

Communications cables are the bane of a PC expert's existence. The problem is that no manufacturer follows the RS232 standard exactly. Many manufacturers use a different connector or they rearrange the order of the pins. Neither is a fatal problem — just a pain in the neck.

Cables can get broken wires and loose pins, or they can be wired incorrectly. The end of this chapter includes wiring diagrams for common PC cables.

10.2.3. The Modem

Modems allow computers to communicate via regular phone lines. They've been around for decades. In the PC world, Hayes Microcomputer Products introduced the Smartmodem (now called the Smartmodem 300) in the early 80's, starting a whole new generation of communications devices.

The Smartmodems are programmable modems. What does "programmable" mean? Well, back in the old days, you had to flip switches to do things like:

- turn echo on/off
- adjust speed, parity, data bits, stop bits
- turn a monitor speaker on/off

Additionally, you had to dial your calls by hand. With a "programmable" modem, you can send it commands via software in the computer. Thus, you could write a program to set some communications parameters, dial a number until it answers, log on to the computer at that number, download a file, and log off all without the operator having to be present. This is truly useful, but it also means that you must learn the modem's terse language in order to get the full benefits of your modem. Or, you must trust that the person who wrote your communications software knew what he/she was doing.

10.2.4. The Communications Software

The best known examples of communications software are:

- Crosstalk, by Microstuf
- Smartcom II, by Hayes
- ProComm, a public domain program

If you have a programmable "smart" modem, as most of us do these days, be sure that your communications software understands the modem's language. Most smart modems claim to use the same language as the market leader, Hayes, but if you're using a clone modem of some kind it can't hurt to check with the software vendor. Most software has been checked out on the more popular clones, like the U.S. Robotics modems.

10.3. MAINTENANCE

These systems are mainly solid state — no moving parts. Modems tend to run hot, as many are designed around one or two fairly dense chips, so don't pile books atop the modem: make sure it's ventilated. I know the ads all show a phone perched on top of a modem, but I'm not sure it's such a good idea. I know for a fact that it's *not* a good idea to stack up Smartmodems. I've had a fair amount of modem trouble: be sure that you know where the warranty information is. As most modems have two year warranties these days, you may need the information.

The usual stuff applies: as they are large connectors with tiny pins, screw in the cables. Don't make the cables longer than they need to be: it increases the noise level in the cable, and if a loop of cable dangles down behind a table it may get caught in something or may be idly kicked. I know of a case where an analyst routinely braced his feet against the cable — just a nervous habit. The cable was even more nervous.

I mentioned power problems in the chapter on power supplies. *The same surge problems can appear on phone lines.* If the line from the phone station to your computer is aboveground at any point, think seriously about a phone line isolator (sometimes called a modem isolator). This sits between the modem and the phone. Another measure of prevention is an RS-232 isolator. It sits between the modem and the PC. I'd rather smoke a modem by itself than a modem and the rest of my system. Examples of such products are offered by the Black Box Catalog:

AC/Modem Surge Protector (contains 2 AC outlets and phone protection)
Stock #GB-SP135
\$139.00

RS-232 In Line Surge Protector
Stock #GB-SP140
\$48.00

The Black Box Corporation
P.O. Box 12800
Pittsburgh, PA 15241
(412) 746-5530

10.4. TROUBLESHOOTING

Communications troubleshooting can provide you with hours of mind-exercising delight, *if* you're the kind of person who likes puzzles. Let's see how to decompose the problem and isolate the bad component.

10.4.1. The Software First

By far, the most common communications problem is the software. You must ensure that both parties in the communication have agreed on the *communication parameters*:

- speed
- parity (can be Even, Odd, or None)
- # data bits (can be 7 or 8)
- # stop bits (generally 1, could be 2)
- local echo on/off

All communications packages have the ability to set these values. Don't worry about what the parameters mean — just be sure that they match. Common symptoms which may be software problems:

Cannot Connect. Check the parameters. If you cannot connect, or you connect and see garbage ("{" characters in particular), you've probably got a speed or data bits mismatch.

Cannot See Input. You type, but you can't see what you're typing. Try a simple command, like whatever command you would use to get a file directory from the distant system. If the command is responded to correctly, the problem is that you must set your local echo ON.

Input is Double. You type a character, and two appear on the screen. Output from the other computer is normal. Answer: set your local echo OFF.

Sometimes, there is no local echo command. In that case, the command has been misnamed (it's a common mistake) to "duplex." Then,

Half Duplex	=	Local Echo ON
Full Duplex	=	Local Echo OFF

Again, it's wrong terminology, but who cares as long as you can get the problem solved.

Cannot Dial. You set up to dial the distant computer, but the modem doesn't respond. There are lots of possibilities:

- Is it plugged into the phone jack? It's common to share with a phone and forget to plug it back in.
- Is the modem plugged into the computer? Did the cable fall out? (It won't if you screw the connectors in.)
- Is the phone line working? Plug a phone in. Is there dial tone?
- Is it your phone system? If you have an in-house phone system, like a ROLM or Northern Telecom system, you may have to issue extra commands to dial out of your system (called a PBX). Check with your telecommunications manager.
- Are you using a high enough quality carrier? If you are communicating at 2400 bps or faster, you may find that the only choice is (no, I don't own stock) AT&T. I know of only one modem (the US Robotics Courier HST) that can handle 2400 bps over a non-AT&T line, and even then its success is sporadic.
- Here's the software one. Is your software set up to talk to your modem? You instruct most modems to dial a number by sending the command ATD, followed by the phone number. Some modems don't do it this way, however. As I said before, is your software equipped to talk to your modem?

10.4.2. The Port

The RS232 port may be ill. As it is a bidirectional device (i.e., it talks as well as listens), you can use a *loopback* to allow it to test itself. Serial ports *do* become ill. I don't know why — I presume it's something to do with the plugging and unplugging that I do with the system that I have (I do a large amount of communications consulting, so I try a lot of equipment out on my system). For instance, I have had two Toshiba 1100+ laptop computers, and they both had something wrong with the serial port: Receive Data in the first one, Ring Indicate in the second. Probably 20% of the expansion boards that I've worked with had a serial problem. So it's worthwhile knowing how to diagnose such a problem.

To test the port, you need to see what is happening on the various control lines. I recommend that you get a *breakout box* to assist you in testing the communications port. A simple, portable, cheap in-line breakout box is available from Hall-Comsec:

WireTap (about \$40)
Hall-Comsec, Inc.
901 Sandy Grove
Fort Collins, CA 80525
(303) 223-8039

Red lights indicate the presence of voltage in excess of +3 volts (interpreted as "0") and green lights indicate the presence of voltage below -3 volts (interpreted as "1"). It's a 2" by 2" mini breakout box that you can carry around conveniently.

10.4.3. Testing The Port

We will use the public domain program PDIAGS.EXE, on the program disk, and a breakout box, like the WireTap. Any breakout box will do, as long as it has indicator lights. You could, alternatively, use a voltmeter, but the breakout box is a good investment.

1. Plug the breakout box directly into the serial connector on the back of the computer. Do not use a cable. You will see activity on just three lines:
 line 2 = -3 volts ("1")
 line 4 = +3 volts ("0")
 line 20 = -3 volts ("1")
2. Activate the diagnostics: PDIAGS. You will see no change in the lights.
3. Get into the serial port diagnostic menu by pressing "S".
4. Open the port: alt-P followed by "O". (Letter "O", not zero.) Line 20 should reverse to value "0". If it does not, your "Data Terminal Ready" (DTR) line is defective.
5. Attach a jumper on the breakout box from line 20 to line 8. On the computer display, the word below "CD" should now be "YES." It was "NO" before. If it does not, your "Carrier Detect" line is defective.
6. Now move the jumper so that it runs from line 20 to line 6. On the computer display, the word below "DSR" should now be "YES." ("CD" will be "NO" once more.) If it does not, your "Data Set Ready" line is defective.
7. Shut off local echo: alt-P followed by "E".
8. Move the jumper so that it connects lines 2 and 3. Now press alt-F. The "fox" message ("The quick brown fox jumped over the lazy dog 0123456789 times!") should appear. If it does not, either your TXD ("Transmit Data") or your RXD ("Receive Data") lines are defective.
9. Quit by pressing alt-Q, then "Q".

If there is a problem, how can you fix it? Obviously, swapping the board is an alternative. But I've fixed a few serial ports when the wires connecting the printed circuit board to the connector were broken, or when the main chip failed.

The "main chip" in most RS-232 interfaces is a large, easily identified chip with an identifying number "8250." It is called a UART, a Universal Asynchronous Receiver/Transmitter. It is generally socketed. They cost about \$4 the last time I bought one. Keep a few around. If the chip is socketed, as it usually is, the swap test takes only minutes.

10.4.4. The Cable

If the port checks out okay, maybe it's the cable. Detach the breakout box from the computer, then connect the cable that you usually use to the computer. Connect the breakout box to the free end of the cable, and re-run the above tests. If one fails, you have a broken wire.

Remember that with cables, speed trades off for distance. If you have a hundred foot RS-232 cable, you won't be able to run it at 38,400 bps. The exact same equipment and a two foot cable might be able to. The only difference is in the length of the cable.

10.4.5. The Modem

Some modems have loopback capabilities built right in. In that case, you can do simple loopback tests. If not, the best first procedure is just a modem swap. They're small and don't require much work to swap.

Do the front lights tell you anything? Put a breakout box in-line between the modem and the computer. You should see the following correspondences:

Modem Light	Red Light On Breakout Box
TR	20
MR	6
RD	3
SD	2
CD	8

When a specific modem light is on, the corresponding breakout box light should be on.

Are the connections correct? Many modems have two phone jacks — one to connect to the wall and one to pass through to the phone. Connect the "phone" one to the wall and nothing happens.

Do the modems match? Are you calling a 9600 bps modem with a 1200 bps modem? If so, is the 9600 smart enough to drop back to 1200? Is it an error-correcting modem, requiring another error-correcting modem in order to carry on a conversation?

10.4.6. The Phone Line

The phone line has four lines, colored red, green, yellow, and black. Only the red and green are used. They should offer 48 volts DC power. If your phone varies greatly from this, call your repair office. Can you try the call on another line? Is it a multi-extension phone — could someone be picking up the phone as your modem tries to dial?

10.4.7. The Other Side

Half of the possibility for problems comes from the folks that you're trying to communicate with. Double check their parameters with them. Try calling out to another computer, if possible. Have they recently installed a new revision of *their* communications software? Installed a new, "completely compatible" modem?

10.4.8. Going Further

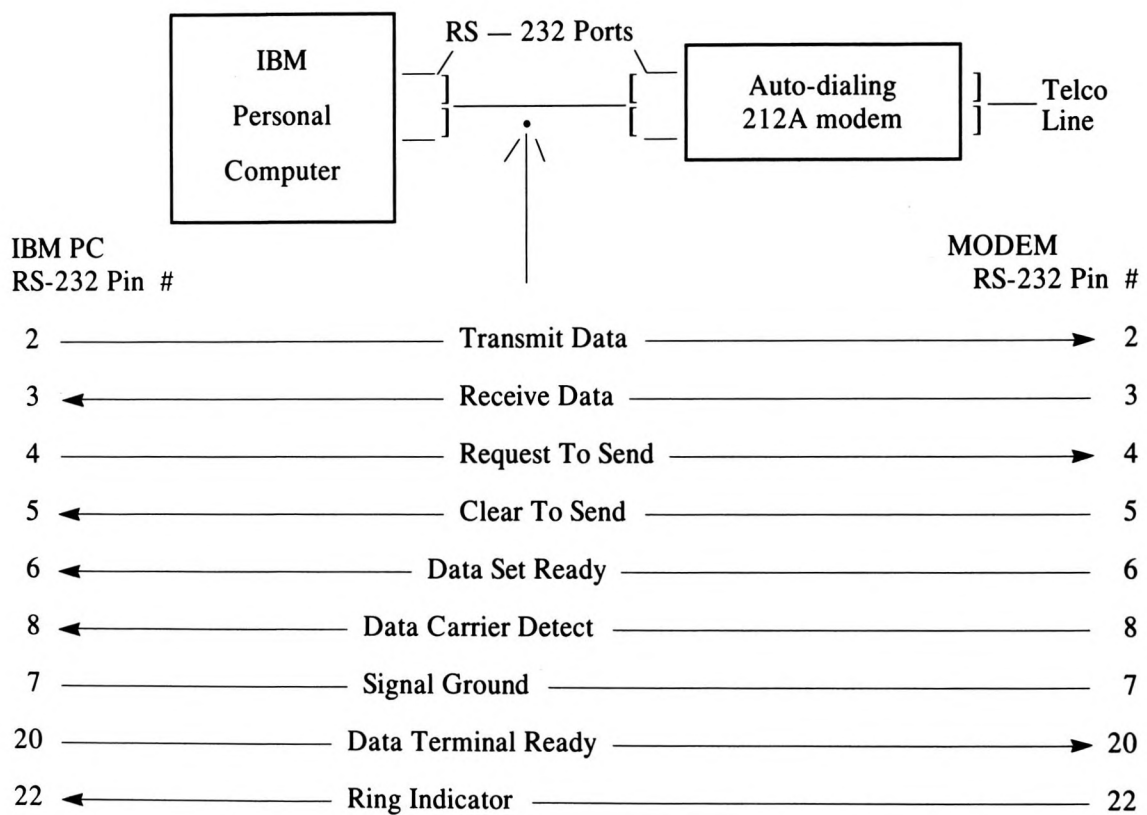
Devices exist, called analog channel test units, which will actually test the phone line between you and the other party. They are, however, quite expensive. An oscilloscope can help you measure the frequencies being produced by your modem to see if they are within specification. Again, you may not want to go this far, as the required investment in test equipment is not trivial.

10.5. COMMON CABLES

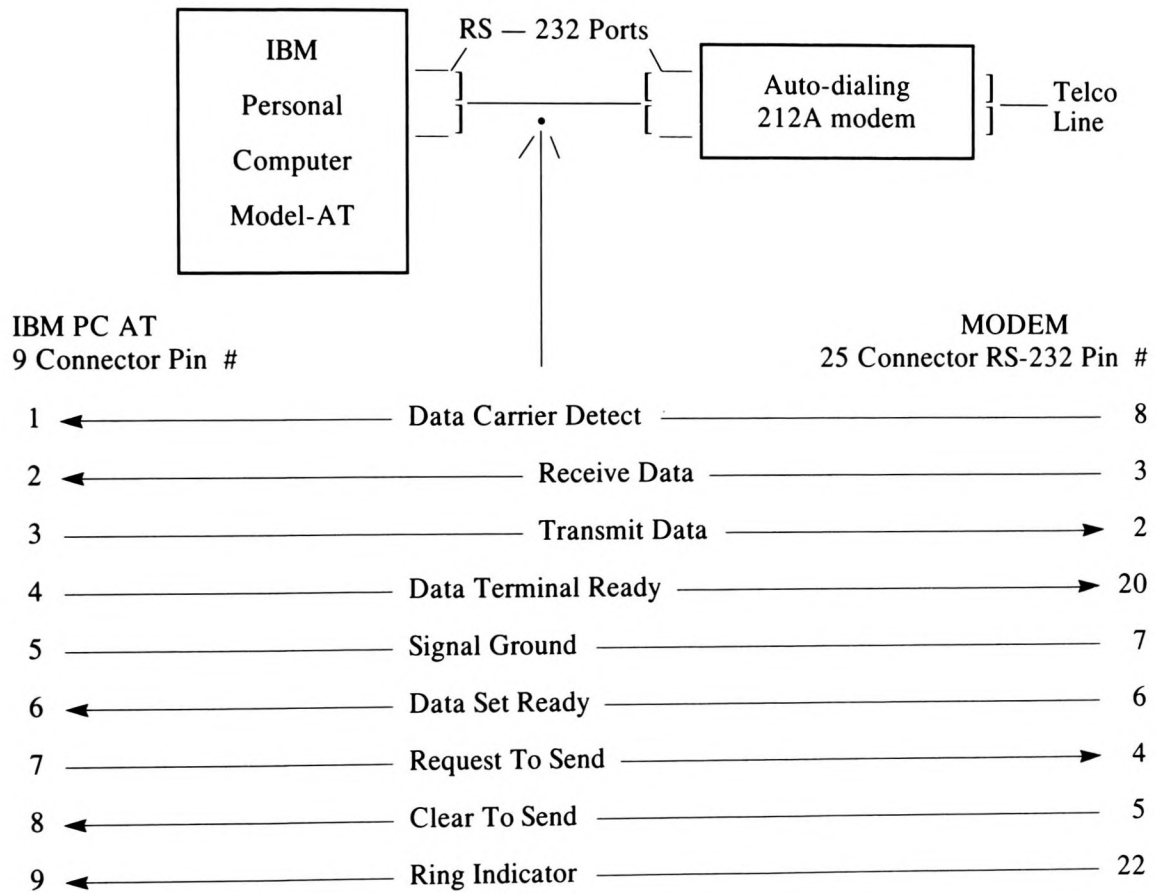
It's not my intention to make an RS-232 expert out of you. But I *can* give you cabling diagrams for some cables that work for some common RS-232 problems. On the following pages, I present cables to connect:

- A PC (25 pin connector) to a modem
- An AT (9 pin connector) to a modem
- A PC (25 pin connector) to another PC
- An AT (9 pin connector) to another AT
- An AT to a PC

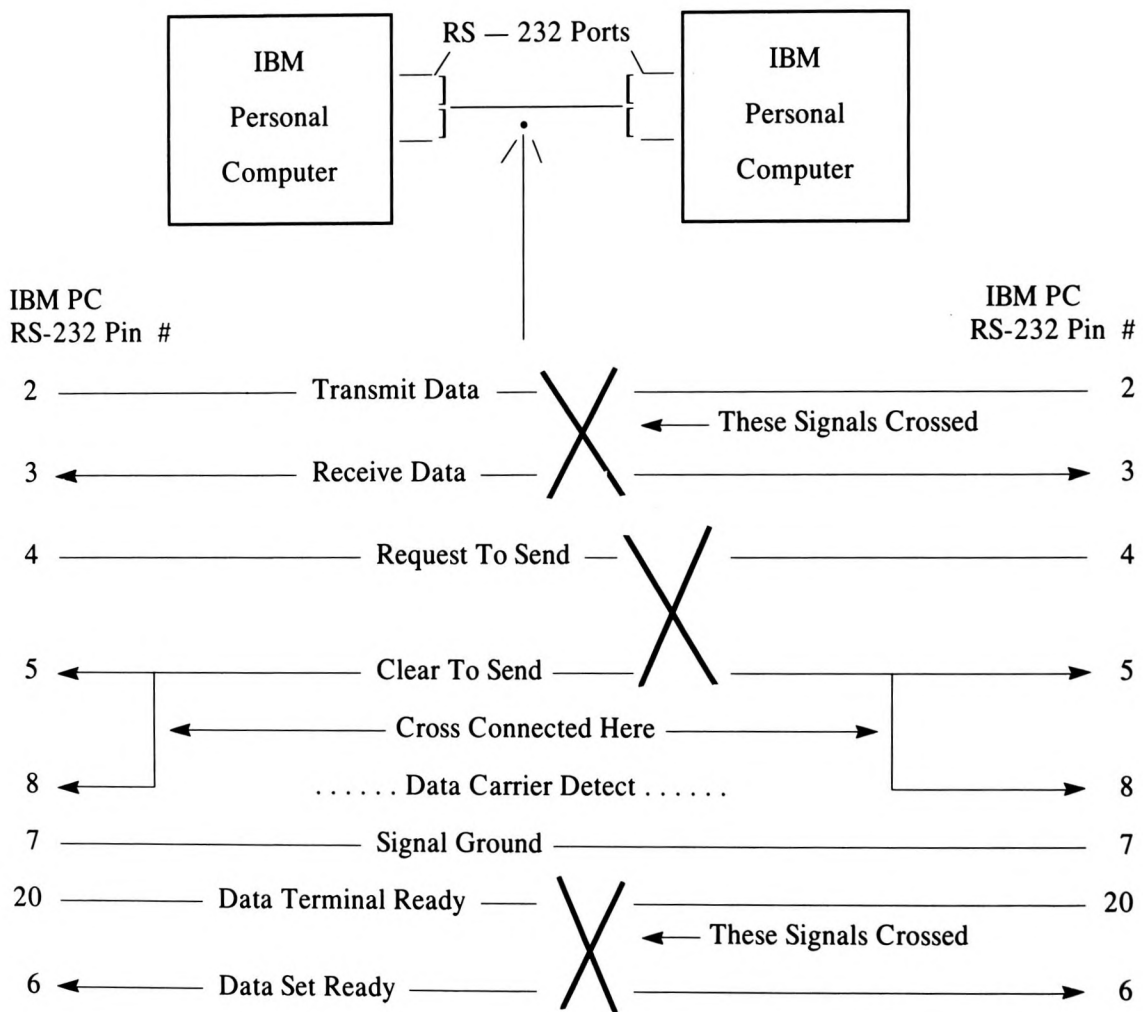
10.5.1 25 Pin PC to 25 Pin Modem



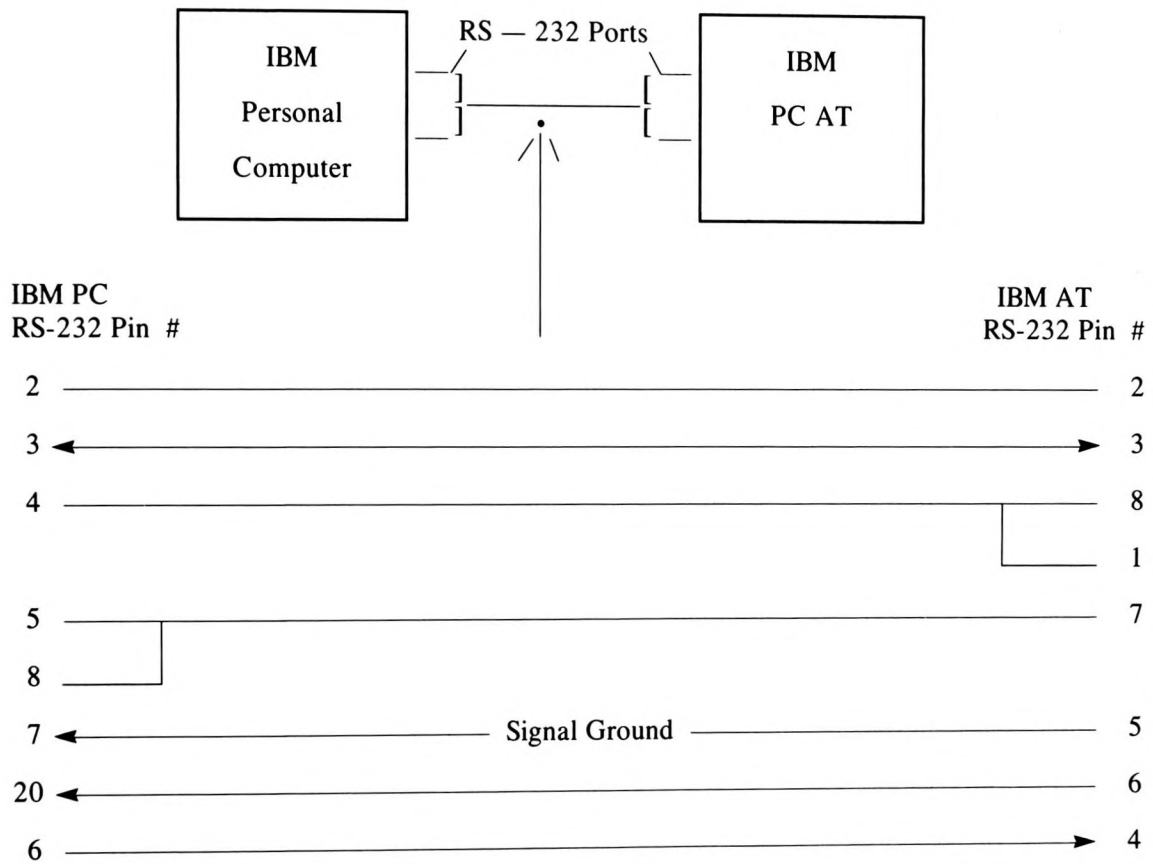
10.5.2. 9 Pin AT to 25 Pin Modem



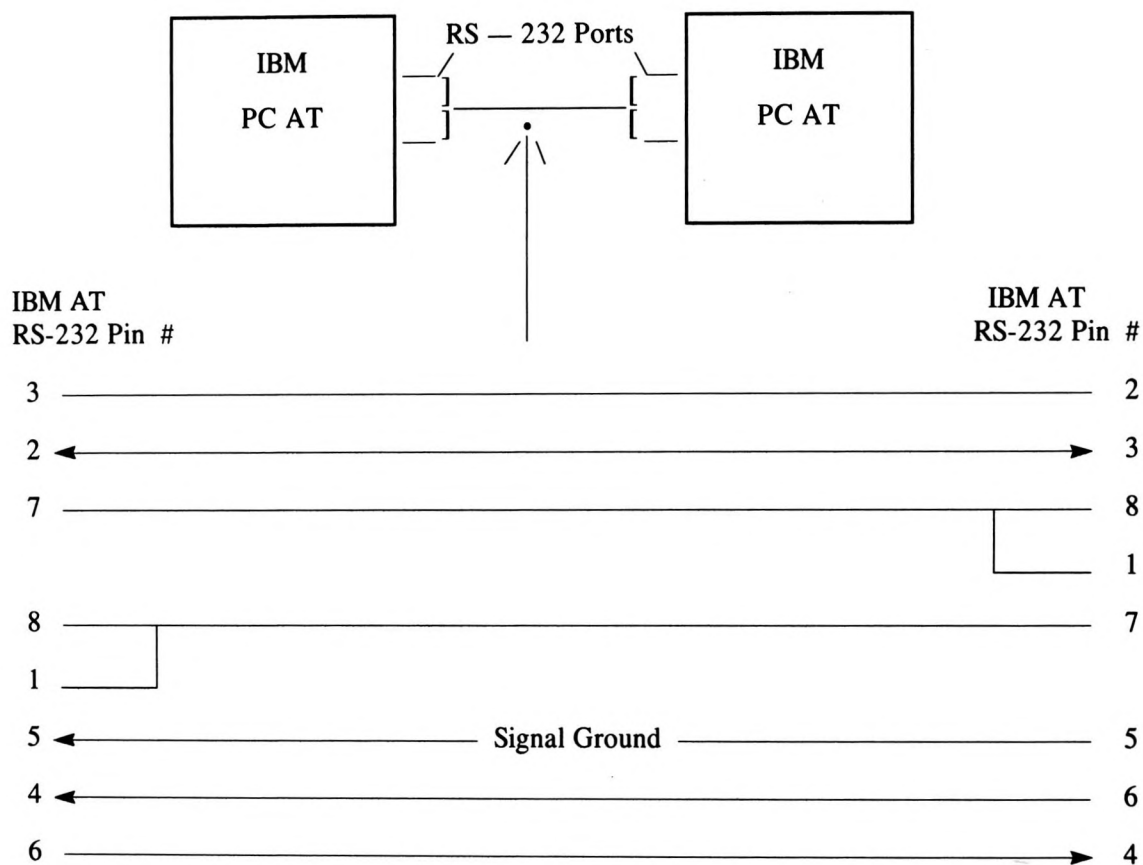
10.5.3. PC to PC Null Modem



10.5.4. PC to AT Null Modem



10.5.5. AT to AT Null Modem



11. KEYBOARDS

The design of the PC keyboard has apparently kept IBM fairly busy in the last few years. Every time a new PC is released, a new keyboard follows. They're all the same to troubleshoot, however.

The PC/XT and AT use different keyboards. If you're swapping keyboards around to troubleshoot a problem, don't swap AT and XT keyboards. XTs and PCs use the same keyboard, however.

11.1. DESIGN AND COMPONENTS

The keyboard is a capacitive one. This is a bit of a departure from the usual. In the old days, a keyboard was just dozens of switches soldered onto a PC board. If a switch died, you'd remove and replace it.


On the PC keyboard, it is one large capacitive surface. You press it at one point and its electronics can sense where that point is. This sounds more complicated (it is), but it involves one large capacitive module rather than a lot of prone-to-fail switches. Unfortunately, it limits you a bit as to what you can do to repair the keyboard.

Each keyboard key pushes a spring, which in turn pushes a paddle. The paddle makes an impression on the capacitive module. The capacitive module sends out signals which are interpreted by the 8048 microprocessor in the keyboard. It, then, sends the key scan code to the PC. The PC then figures out what the key means.

As before, let's see what we can take apart and test.

11.1.1. The System Board Keyboard Interface

First is the interface between the keyboard cable and the system unit. This is a DIN plug with five pins, all numbered oddly. The diagram below is of the connector *on the system unit*.



Test Voltages: voltage between all pins and pin 4 should be in the range of 2 - 5.5 volts DC. If any of these voltages are wrong, the problem probably lies in the PC — the system board in particular. If they're okay, the problem is probably in the keyboard. Check the keyboard cable next.

11.1.2. The Keyboard Connector

The keyboard cable runs from a DIN connector which attaches to the system unit to a flat-jaw type connector inside the keyboard housing. The cable has five wires, and can be checked for continuity quickly with an ohmmeter: disconnect the cable inside the keyboard and run the tests. This is discussed at greater length later.

11.2. MAINTENANCE

The major item of maintenance for keyboards is *abstinence*. Abstinence, that is, from spilling things into the keyboard. The Saf-T-Skin was discussed in the Maintenance chapter: this is one protection approach. Another is to just be careful.

Periodically disconnect the keyboard and remove it from the general area of the PC. Then hold it upside down and shake a bit to remove the small debris which lives there. Blow out the keys with compressed air.

11.3. TROUBLESHOOTING

As with other items, you may want to view the keyboard as disposable. Replacement keyboards cost about \$69 from discount houses. There are some simple things which you can do before throwing out a keyboard, however.

11.3.1. Is it Plugged In?

On the back of PCs and some compatibles there are two identical ports, a cassette port and the keyboard port. Plugging the keyboard into the cassette port won't work. You could even destroy the keyboard if cassette instructions are sent to the cassette port. The ports should be labelled, but if they're not, the keyboard port is the one closer to the power supply.

11.3.2. Is It One Key, Or All Keys?

If only one key is malfunctioning, check that key's spring. Remove the key by grabbing it with your fingers and pulling up. For the tough keys, fashion a hook from a paper clip. Under the key you will see a spring. Replace the key cap and see if the problem goes away. If not, try pulling the spring out *just a little*. Then replace the key cap.

11.3.3. Checking The Cable Continuity

Next, test continuity of the cable. Turn the keyboard upside down so that the cable is coming out of the back of the keyboard, to the right. Remove the two screws. The bottom plate will swing back and up to remove.

You will now see that the cable splits to a single wire, which is grounded to the bottom plate, and a cable with a flat-jaw connector. Push apart the jaws of the connector to release. You can then test each of the five wires.

11.3.4. Completely Disassembling The Keyboard

Not recommended for the fainthearted. You've got a good chance of making things even worse, so don't do it unless there's no other hope.

Remove the main assembly from the keyboard case. With a vise, or C clamps, set up a support for the keyboard assembly on the sides. If you don't support it on the sides, all Hell will break loose when you remove the back metal plate.

A printed circuit board with capacitive pads on it is held against a metal plate by ten metal tabs — five above, five below — and a hex screw. You must use a set of pliers to unbend the tabs enough to remove the plate. Position the assembly so that the plate is on top (not the printed circuit board), then remove the metal plate.

The plate has been holding dozens of plastic paddles against the printed circuit card, one for each key. These are small, flat, easily-broken pieces of plastic shaped like the outline of a castle turret, or the "rook" symbol in Chess. What you are looking for is a broken paddle. Replace any broken paddles and replace the metal plate.

The other reason to attempt this is if you've poured hot coffee into the keyboard. Very carefully remove all of the paddles and clean them as well as the bracket that they sit in. Then use alcohol to clean the capacitive PC board. Then reassemble and pray. I have seen this work, but only once.

Replacing the metal plate is a bit difficult. Position it correctly, then use clamps to hold one side together while you bend the tabs on the other side into place. Again, it is very easy to damage paddles at this point. Where do you get replacement paddles? No one sells them, that I know of. Get them from the first keyboard that dies. It's good for 83 replacement paddles.

11.4. REPLACEMENT KEYBOARDS

If you've gotten this far and no luck, don't despair. Many keyboard problems are not soluble, and the things aren't that expensive anyway.

Keytronics and Datadesk make good replacement keyboards. Datadesk's are a bit cheaper, and they have one that is programmable. What that means is that you can move that annoying ESC key somewhere else, set up your keyboard as a Dvorak keyboard, and the like.

Keytronics offers a keyboard with in-line bar code reader. Other software allows you to print out bar codes with a dot matrix or laser printer. One firm I know of uses bar codes at the top of their document tracking forms so that secretaries don't have to key (and sometimes miskey) in a twelve digit document tracking ID every time they have to update the document tracking database.

Another interesting keyboard enhancer product is a touchpad. Keytronics offers the KB5153, which can be AT or XT compatible and includes a capacitive touchpad. It is considerably more expensive — about \$280 — but that, recall, is the price that IBM charged for the original simple PC keyboards.

12. DISPLAYS AND DISPLAY ADAPTERS

The display on a computer is the primary output device. It's the kind of thing that people get opinionated about. Personally, I've had the cheapest video combination available for years — simple composite black & white monitor and color/graphics adapter — but some folks just gotta have high resolution displays. What's behind all this, and how much can we fix?

The display system has two components: a display and a display controller.

12.1. COMPONENT: DISPLAYS

The first part of a display system is the display. Many options have grown up in the PC world. Some are cheap, some cost thousands of dollars. They all show at least 25 lines of 80 column text.

There are several basic kinds of monitors:

- Composite monochrome (around \$75)
- Composite color (\$200 - \$350)
- RGB (\$350 +)
- EGA monitors
- Direct drive TTL monochrome (\$150)
- Multiscan (\$500+)

Composite monitors are a step better than a television set. They offer greater resolution than a TV, but that's about it. They offer low expense and good quality.

RGB ("Red, Green, Blue") monitors offer better resolution for color text than do color composite monitors. The IBM Color Display and the Princeton Graphics HX12 are RGB monitors. RGBs are more expensive, offer somewhat better quality than composite color, and attach to a display adapter with a nine pin D shell connector.

More expensive than regular RGB, EGA (Enhanced Graphics Adapter) monitors are RGB monitors with higher resolution. Characters are sharper and the display steadier. In addition, these monitors can display 43 lines on a screen rather than the usual 25. This isn't as great as it sounds: most programs won't support the 43 line mode. Notable exceptions are dBase and the IBM Personal Editor. An example of this type of monitor is the IBM Enhanced Color Display.

Direct drive TTL monitors are monochromatic (generally amber and black or green and black) and offer high resolution, steady images. They are connected with a nine pin D shell connector like the RGB, but don't plug one of these in where RGB is expected or vice versa. An example of this is the IBM Monochrome Display. You honestly will get smoke.

Fairly expensive, but most versatile, are the multiscan monitors. These can serve as an RGB or an EGA monitor — they sense at what rate data is coming in, and adjust their scan rates accordingly. The best-known example is the NEC Multisync monitor.

12.2. COMPONENT: DISPLAY ADAPTER BOARDS

The second part of a display subsystem is the controller board. There are quite a few in use currently, and they all have their characteristics. A table, following, summarizes them.

Common PC Adapter Types

Adapter	Maximum Resolution/ # Colors	Maximum Character Box	Monitors Supported	Graphics Supported?
Color/Graphics	640 x 200/2	8 x 8	RGB Composite	Yes
Monochrome Display	720 x 350/2	9 x 14	Mono TTL	No
Enhanced Graphics	640 x 350/64	8 x 14	RGB EGA Mono TTL Composite	Yes
Hercules Graphics	720 x 350/2	9 x 14	Mono TTL	Yes
Multi Color Graphics Array (MCGA)	640 x 480/2	8 x 16	Hi Scan RGB	Yes
Video Graphics Array (VGA)	640 x 480/16	9 x 16	Hi Scan RGB	Yes
8514/A (an add-on adapter for VGA computers)	1024 x 768/16	9 x 16	Hi Scan RGB	Yes

12.3. UNDERSTANDING THE TERMINOLOGY

As you probably know, CRTs work by painting lines of dots on the screen. At any point in time, only one dot on the whole screen is turned on. You don't see this, fortunately, as your eyes contain *visual purple*, which is responsible for *persistence of vision*. To maintain the illusion of a steady screen, about 60 screens must be painted onto the CRT each second. Significantly less screens will lead to flicker. Some people have very fast eyes, and see flicker in normal screens. I am thankful for slow eyes.

The rate of screens per second is called the *vertical trace frequency*. It is 60 hz for most screens, 50 for monochrome TTL monitors. The TTL monitors often combat the flicker problem by using high persistence phosphor on the inside of the CRT. This adds persistence, aiding your eye. This is not expressed in screens/second, but hertz. Not correct, but everybody does it.

The vertical scan frequency is more interesting. The monitor must draw several hundred lines in order to display one screen. As said before, it must also display around 60 screens per second. A composite monitor, for instance, uses 262 lines per screen. You can only use 200 — the rest are wasted. 262 lines/screen times 60 screens/second equal roughly 15,700 lines per second. This is referred to as a 15.7 Khz *vertical scan frequency*. Yes, hertz is really not a good unit, but it's the one everyone uses. The table below has the values for common monitor types.

Statistics for Common Monitor Types

Type	Screens/Second	Lines/Screen		Vertical Scan Frequency
		Actual	Used	
Composite	60	262	200	15.7 Khz
RGB	60	262	200	15.7 Khz
TTL Direct Drive	50	368	350	18.4 Khz
EGA	60	363	350	21.8 Khz
Multiscan Monitors	depends	up to 500		up to 35 Khz

Another monitor buzzword is *dot pitch*. This is how narrow each dot is. The narrower, the better: the characters can be sharper. Nowadays many manufacturers offer .31 mm or better.

12.3.1. Interlacing

Some computers, like the Commodore Amiga, seem to get unbelievable resolution out of regular composite monitors. One PC board, the EGA Wonder from Array Technologies, Inc. (they're in Ontario at (416) 477-8804), shows EGA graphics on a regular composite monitor.

How is this possible? An EGA display needs 350 lines, and the absolute maximum that a composite has is 262 lines.

The trick is this: it actually shows each screen as *two* screens, one immediately following another. The even lines, then the odd lines. There's no free lunch, however: this means that instead of 60 screens, we get 30 screens/second, as each "screen" requires two screen paints. This process of alternating odd/even to get greater resolution is called *interlacing*. The screen uses the top 175 lines to display 350 lines. Thus, it not only looks weird, it looks squat and weird. It makes the screen look flickery. Adding a glare screen reduces this somewhat.

12.4. MAINTENANCE

Not much to say here. Treat the display board like another PC board. Ensure that the connectors are secured, as the little pins can bend over time. Be very sure not to plug RGB into monochrome and vice versa, even if they *do* use the same connector.

Clean the display when it gets too dust-filmed. One item to be careful about, however: I have heard of a case wherein a NEC Multisync monitor actually imploded when cleaned with a damp cloth while turned on. I'm not sure that I believe it, but maybe you want to turn the monitor off before cleaning it.

12.5. TROUBLESHOOTING

Some of the dumbest monitor problems are the easiest.

- Is it turned on?
- Is the brightness or contrast turned down?
- Is everything plugged in? Is it plugged into the right place?
- If you are using a multiple display board, are the DIP switches and jumpers set for a mono display or RGB? If you set it for RGB and plugged in mono, throw away the mono monitor.
- Has someone cleverly convinced DOS to display black letters on a black background? Reboot.
- Did you hear one long and two short beeps indicating a bad video card?
- If using a multispeed computer, is it in Turbo mode? The display memory may not be fast enough in the higher speed mode. Drop back to the lower speed. If the problem goes away, you'll probably have to replace the memory on the display board.
- There are non-video reasons for a display "malfunction." Like when the power supply has killed the computer. If the display is dead, do you hear the power supply fan? Try typing DIR blind. If the computer is okay but the display is bad, you will see the drive light come on. Use a sound-emitting program to see if the computer is functioning.
- Are the motherboard DIP switches set correctly for your display adapter?

The quickest test is a monitor swap. **DO NOT DO THIS** if you're working with a multidisplay (one which supports both RGB and mono TTL monitors) and are unsure about jumper settings. There's no sense in torpedoing *two* monitors today.

If that does nothing, swap the display cards. Then the cable. If the display card is the problem, do the usual easy stuff: check the seating of the socketed chips. Clean the edge connector and the video connectors. Try again.

Still no luck? If you're courageous, you can attack the display board. Many display boards are based on a Motorola 6845 chip, a character generator ROM, and memory. If the memory is faulty, you'll only lose some of the display. If the ROM is bad, you'll get funny looking characters on screen. If the 6845 is bad, just about anything could happen.

Don't try to service the monitor. As I've said before, you can hurt yourself doing that.

13. THE MOST COMMON PROBLEM: SOFTWARE

Software is probably the most common PC problem. You can't usually "troubleshoot" software in the sense that you can hardware — it doesn't wear out a print head or need replacement after 10,000 uses. You *can* troubleshoot software in another way: determine what is causing the problem. This leads you to the right source for a solution.

There are basically two kinds of software problems, ones which are *your* problem and ones which are the *software vendor's* problem.

13.1. TRACKING DOWN THE PROBLEM

Here's a few tips on solving a problem. Sometimes (hopefully), problems just go away. If they don't, we at least try to fence them in, understanding what makes the bugs come out and what keeps them quiet.

13.1.1. If The Software Has Worked Before

1. Reboot. Sometimes, cosmic rays (I know, it's a weak excuse) or a power surge makes the computer temporarily brain-damaged.
2. If the problem persists, perhaps the program has been damaged on disk. Minor damage — a byte or so — would not incapacitate the program, just produce strange results in some cases. Reload the software from the backups. (You do back up your software, don't you?)
3. However, do not try to restore regular backups of copy protected software unless you have used a backup program made for copy protected software, like Copywrite or Copy II PC.
4. Has someone helpfully messed with your program? Many programs include an Install program to set information like screen type, printer type, etc. Re-run it, just in case. You can generally see if someone has changed the settings by checking the date on the configuration or setup file.
5. Have you started using a new memory resident program, like Metro or Sidekick? Re-try without the new software. In fact...
6. ... re-boot without *any* memory resident programs. I only know of a few that have never given me any conflict problems.
7. Did you install another piece of software recently? If so, did it "helpfully" re-write your CONFIG.SYS or AUTOEXEC.BAT files? I have started making a habit of backing up these two files before installing new software.
8. If you suspect memory residents, do not despair: you may be able to solve the problem by re-arranging the order in which the programs are called.

13.1.2. If You Are Installing Software Now

If you've *never* seen this software work before, there are several possibilities.

1. First, of course, is the PC compatibility question. If you are not using a 100% IBM PC, there is the possibility that your compatible just plain can't run this program. Run it on an IBM PC with the same configuration (if you don't have one, how about the local computer store?). If it runs, that's the problem. Contact your computer manufacturer. Just in case, call the software folks and ask if they've, perchance, found a patch for your computer. If you have a clone, ask if there is a newer version of the BIOS. If so, get the latest ROM.
2. Check the version of DOS. Don't even think of using a version of DOS before 2.1. Use 3.1 or 3.2 preferably. Some programs run only on 2.1 or only 3.1 or 3.2. Obviously, network applications must run on 3.1, 3.2, or 3.3. The occasional program needs so much memory that it can't run in 640K unless you're using DOS 2.1, which is smaller than later versions. Another problem is running a secondary command processor from inside an application. DOS 2.1 doesn't do it particularly well — try 3.X.
3. Check memory. Some programs, oddly, won't run in *more* than 512K. You actually have to run some memory-resident programs to cut down enough memory so the program will run. One answer is to re-set the DIP switches or run the SETUP program so that the computer thinks it has less than 640K. Alternatively, there is a program MEMSET.COM which will do this for you.
4. What is required in the CONFIG.SYS and AUTOEXEC.BAT files? Check the installation notes. What must the FILES= and BUFFERS= values be?
5. Some (thankfully, less and less) programs assume the existence of a B: floppy, and can only read data from there. Or they may require a subdirectory that you have not created. Some programs will warn you of the problem. Others just offer an error number or a cryptic message, then drop you back to DOS.
6. Other programs start and then seem to "hang." A couple of these have fooled me until I realized that they are looking for a particular device, like a printer or a modem, and that they are trying to initialize that device. If it isn't around, they go out to lunch and stay there.

13.2. REPORTING THE PROBLEM

Once you've isolated the problem, report it to the software manufacturer. In all probability, they'll be fairly skeptical. You'll have to prove that the problem exists. It's unfortunate that some vendors charge you money for buggy software, then treat you like a deadbeat when you complain that the software doesn't do what it is advertised to do. In any case, report it to your local user group through whatever "user to user" type column the group has in its monthly newsletter. (You *are* a member of your local user group, aren't you?)

Once you have determined the problem and what causes it to occur, produce the briefest possible example of it. Write up the problem and any workarounds in a letter to the software vendor. Include the brief example on a floppy disk, and send the floppy with the letter to the vendor. Other information to include:

- version number of the software
- the kind of PC you are using

- if it is a multi-speed PC, the speed that you were working at
- the version of DOS that you use

When you call with a problem, vendors will sometimes ask,

"What version are you working with?"

You reply (let us say), "Version 2.3."

"Oh, ho," they respond. "You're *way* behind. We're up to 3.1" (I sent in my registration cards. How is it that I was so uninformed?) "What's your name? I'll send you the upgrade."

At this point, you think that the answer is in the mail. In my experience, new versions fix about only 50% of the problems. Ask, "does the new version fix this problem? Have you ever heard the problem before?"

14. ADVANCED PROJECTS: SAFELY IMPROVING YOUR PC OR AT

By now, you have the flush of power. No longer is the PC the daunting, mysterious device it once was. So you might be feeling brave. Here's a few parting projects: cheap improvements on the PC.

14.1. SPEAKER VOLUME CONTROL

One thing on the PC that could have been done a bit better. The speaker. It'd be nice to turn it down, or shut it off, or connect the output to another speaker.

The speaker is connected to the PC via a flat rectangular connector called a *Berg connector*. It's got four connectors, but only two are used. It's a simple circuit: two wires attach to the leads of the speaker. The speaker is an 8 ohm speaker.

Some people disconnect the speaker to shut it up. This isn't a good idea, from what some engineers tell me. Apparently you need to replace the speaker resistance with some alternative of the same resistance, like a resistor. The circuit is a 1/2 watt circuit, so don't use the smallest resistor available: get a 1/2 watt power resistor. Then you want a DPDT (Double Pole, Double Throw) switch so that one position connects the speaker wires to the speaker, and the other position connects the speaker wires to the resistor.

A superior approach would be to put a variable resistor (called a *potentiometer*) in line with the speaker. The speaker would always be in the circuit, but the potentiometer could produce enough resistance to make the speaker quiet or inaudible. You could run longer wires from the pot to the speaker so that the pot could serve as a volume control. A potentiometer with a maximum resistance of 100 to 300 ohms would serve well. (Minimum resistance on potentiometers is zero ohms.)

14.2. An Interface For An External Speaker

It may be that you use your speaker for something more than beeps, like tinny music or awful speech synthesis. In this case, you might want an external speaker jack so that you could run it into your stereo, or into a recording device.

You would use a typical phone jack, but a better approach would be to use a two-position jack. When nothing is plugged into the jack, the internal speaker would be in use. When something is plugged in, the internal speaker is taken out of the circuit and the external is in the circuit.

Before plugging this into anything, be sure that it is designed for 1/2 watt power, 8 ohm resistance.

14.3. Installing A Hardware Reset Switch

Something that many clones have that the PC does not have is the hardware reset switch. The ctrl-alt-del combination is a software combination: if the processor meets a simple HLT (halt) opcode, no key combination will be of any use.

The chip that controls RESET is the 8284, the clock chip. Just hold its RESET line to ground (you must be sure to hold it down for at least 810 nanoseconds), and the process is done. RESET is pin 11, GND is pin 9. Just install a pushbutton switch between 9 and 11. Put a 300 ohm resistor in line with the switch so as to not cook the chip. And that's all there is to it.

15. APPENDICES

15.1. THE BASIC SURVIVAL KIT

What constitutes a basic troubleshooting tool kit? I have mentioned the following tools at various times:

- Hex nut drivers
- Tweezers or hemostats
- Multimeter
- Digital Temperature Probe
- Compressed Air
- Tweaker screwdriver or chip puller
- Software:
 - The Program Disk (Public Domain Utilities)
 - Mace Utilities
 - The Mace Advanced Hard Disk Utilities (Kolod Utilities)
 - The Norton Utilities

15.2. THE PROGRAM DISK

I have collected a number of public domain programs that can be of help to you in maintaining your computer. Please understand that these programs are public domain or shareware: we are NOT selling these programs, nor do we offer any license or guarantees about their performance. Please use them with caution.

Some programs are referred to in the text. Some are not. All programs either come with a file with extension .DOC or .TXT or the like, which document the file. You should not try to use the programs until you have printed out and read these files. Some programs do not come with files, but rather will show on-screen documentation when invoked with "?", as in CORETEST:

CORETEST?

will show a few screens of documentation on CORETEST.

I can't stress how important it is to understand what a program is trying to do before attempting to use the program. For example, you can irrevocably format your hard disk with the HDAT program.

In order to put over 450K on a single disk, the files have been squeezed with a file compression utility. The utility is included on the disk. If you have a hard disk, the batch file INSTALL.BAT will unsqueeze all of the files for you. Just put the program disk in drive A: and type

A:INSTALL

and the programs will be installed in a subdirectory on your hard disk.

The files on the disk are as follows:

INDEX	A listing of all file names and descriptions.
PDIAGS	Port Diagnostics.
NMI	Intercepts parity errors. Invoke: NMI/I

DISKPARK	Various drive park programs.
AUTOPARK	Automatically parks the drive after a certain amount of time.
LPT2DSK	Redirects printed output to disk.
DSKMON	Informs you of disk errors that DOS doesn't tell you about.
INTERLEV	Computes optimal hard disk interleave.
HDAT	Good set of hard disk and controller tests.
CDISK	Lets you write protect drives, even fixed ones, and logs access times.
EQUIPCFG	Gives status info, ROM dates.
DIPSET	Aids in setting DIP switches on PC or XT motherboard.
INFO	Gives disk info.
HDTABLE	Shows you your AT's disk parameter table.
CORETEST	Tests disk performance - did you get what you paid for?
SPINSPED	Floppy RPM speed tester.
MEMSET	Set system memory to anything you like.
ARC521	Used to unsqueeze the files.